DIPARTIMENTO DI MATEMATICA - UNIVERSITÀ DEGLI STUDI DI GENOVA

Dottorato in Matematica ed Applicazioni - Ciclo XXXI

Two Studies in Representation of Signals

Ph.D. thesis of

Enrico Cecini

Supervisor

Prof. Ernesto De Vito

To be defended in May 2020 at the Department of Mathematics of the University of Genova

Referees: Ilaria Giulini, Emanuele Delucchi

Comments or suggestions are welcome, please contact cecini@dima.unige.it

This thesis consists of two separate parts:

The first part, Chapter 1, is based on:

E. Cecini, E. De Vito, L. Rosasco, *Multi-Scale Vector Quantization with Reconstruction Trees,* to appear in - *Information and Inference: A Journal of the IMA*, Oxford University Press.

The second part, Chapter 2, is based on:

E. Cecini, *Matroidal Structures in Graph Signal-Processing,* manuscript.

Preface

The common denominator of this thesis is, loosely speaking, the treatment of signals towards a compact and meaningful representation, in particular taking advantage of mathematical tools of discrete nature; the concrete meaning that these terms take varies across the two parts of the thesis, and is properly defined therein. The exposition of each part is self-contained and with no cross-reference; extensive review of the literature is offered for each topic separately.

Common keywords of the two parts are probably: subsampling, partitions, partial orders, graphs (including trees), wavelets.

The two parts are presented in chronological order. Differences in style of exposition reflect the different stage of development of the two projects (by the delivery of this thesis) as well as the different common practices of each subject matter.

Despite the common ground, the mathematical flavour of the two parts is very different. The drift from one subject to the other, even though seemingly radical, was not the result of a plan, but happened gradually and organically. Subsampling of point-clouds is the main subject of the first project. On the one hand, constructing a graph from a point-cloud is the first step of several algorithms reviewed at the time of the first project. This triggered an interest for Graph Signal Processing. On the other hand, while considering several subsampling strategies to compare with, the attention fell onto Determinantal Point Processes. The study of the latter resulted in a growing interest towards an algebraic foundation of the theory of Graph Signal Processing, that led to the use of Matroid Theory. The complexity of subject outgrew the expectations, pushing out of reach some of the former targets. In particular Probability Theory disappeared from the scope of the second project while Combinatorics was taking over.

The first part is much closer to a direct application, while the second is probably more theoretically ambitious.

Contents

CHAPTER 1

Multi-Scale Vector Quantization with Reconstruction Trees

1. Problem Statement

The setting at hand is of interest from various points of view. For this reason we would like to provide the reader with a sketch of the mathematical problem itself, before contextualizing it with motivations, background and connections with related lines of research.

A random vector X is given, taking values in a bounded region $\mathcal{X} \subset \mathbb{R}^D$, and distributed according to a probability measure ρ, so that $\text{supp}(\rho) \subseteq \mathcal{X}$. We consider n identical and independent copies of X, denoted by $X_1, \dots, X_n$. A sample $x_1, \dots, x_n$ drawn from the above random variables will be called *dataset*, possibly arranged as empirical measure $\widehat{\rho} = \sum_{i=1}^n \delta_{x_i}$. The dataset is meant to be the only input, while ρ remains unknown.

The aim is determining a set of $N \leq n$ vectors in $\mathcal{X}$, that will be called *centroids*, $c_1, \dots, c_N$, together with a function P sending each point of $\mathcal{X}$ to a centroid. To this end, each centroid c_i is associated with a subset $I_i \subseteq \mathcal{X}$, called cell. These cells are such that $\Lambda = \{I_1, \dots, I_N\}$ defines a partition of $\mathcal{X}$; namely $\bigcup_{I \in \Lambda} I \supseteq \mathcal{X}$ and $I \cap J = \varnothing$ whenever $I \neq J$. Then, the map $P_\Lambda : \mathcal{X} \to \mathcal{X}$ is defined by

$$P_\Lambda(x) = \sum_{j=1}^{N} c_j \mathbb{1}_{I_j}(x), \tag{1}$$

for $x \in \mathcal{X}$; therefore sending each point of I_j to c_j. With slight abuse of notation we call P_Λ a nonlinear projection.

The following functional will be called *error*, or *distortion*, induced by P_Λ:

$$\mathcal{E}[P_\Lambda] = \mathbb{E}[\|X - P_\Lambda(X)\|^2] = \int_{\mathcal{X}} d\rho(x) \, \|x - P_\Lambda(x)\|^2 , \tag{2}$$

where $\|\cdot\|$ is the Euclidean norm in $\mathbb{R}^D$. We aim at providing a computationally cheap way to determine a projection $\widehat{P}_\Lambda$ depending only on the data $x_1, \dots, x_n$, in such a way that it has low expected distortion.

Besides computational convenience, the choice of N (hence the design of Λ) is desired to be flexible. This in turn is related with the idea of a *coarse-to-fine* representation of the dataset. Let's introduce a *scale* parameter $\eta \in \mathbb{R}^+$, used to gauge the complexity of the partition Λ. A coarse-to-fine family of partitions Λ_η is such that:

$$\eta_1 > \eta_2 \quad > \dots \tag{3}$$

$$\Lambda_{\eta_1} \supseteq \Lambda_{\eta_2} \supseteq \dots \tag{4}$$

$$N_{\eta_1} \leq N_{\eta_2} \leq \dots \tag{5}$$

where the notation $\Lambda_{\eta_1} \supseteq \Lambda_{\eta_2}$ means that for all $I \in \Lambda_{\eta_2}$ there exists $J \in \Lambda_{\eta_2}$ such that $I \subseteq J$, that is, Λ_{η_2} is a *refinement* of Λ_{η_1}. In other words, each cell of Λ_{η_1} gets split in a number of cells in Λ_{η_2} (or remains unchanged).

The operative definition of Λ_η (depending on the data $x_1, \ldots, x_n$) is deferred to later sections, but, regardless it, $\mathcal{E}[P_{\Lambda_\eta}]$ would in general decrease together with η. Our main contribution consists in an estimate of the form:

$$\mathcal{E}[P_{\Lambda_\eta}] \leq \epsilon(n, \eta) \tag{6}$$

under suitable assumptions, therefore providing statistical guaranties on the accuracy of using $P_\Lambda(x)$ as an approximation of x.

2. Background

Dealing with large high-dimensional data-sets is a hallmark of modern signal processing and machine learning. In this context, finding parsimonious representation from unlabelled data is often key to both reliable estimation and efficient computations, and more generally for exploratory data analysis.

Indeed the problem of dealing with large high-dimensional data is twofold. On the one hand complex algorithms simply cannot run in real time on such large datasets. On the other hand, besides the practical inconvenience, the analysis may qualitatively benefit from some sort of compressive preprocessing, or, as it shall be referred to, a parsimonious representation. This idea is based on the assumption that the relevant information carried by the data is less voluminous than what size and dimensionality may suggest. A parsimonious representation is supposed to highlight relevant information and put aside redundancies. Indeed here comes into play the other key concept motivating our work, that is the multi-scale approach to data representation: even if there is no redundancy, the information can still be filtered according to some measure of relevance, in order to be processed accordingly, rather than in an unstructured batch.

2.1. Dimensionality and Representation.

Even being distinct features, size and dimensionality of a dataset are in fact related once a certain task to be carried out is fixed, as effective dimensionality affects the minimum size necessary to perform that task.

The term *Course of dimensionality* was coined first by R.E. Bellman in the late 1950s to describe the typical difficulties occurring whenever low dimensional intuition is naively carried over high-dimensional problems [8, 31, 52].

Referring to the problem sketched through the previous section, the term *dimesionality* here broadly refers to various definitions of dimension that one might associate to the support $S = \mathrm{supp}(\rho) \subset \mathcal{X}$ of the probability measure of the data ρ. Most definitions are metric in nature, and deal with the idea of covering or packing the support with open metric balls (D-dimensional in our context). We briefly review the definitions of Minkowski-Bouligand dimension and Hausdorff; even if they are not used explicitly in our results, they do provide useful insights about our main assumptions. The Minkowski-Bouligand dimension [37] is based on relating the radius r of the balls with the number of balls needed to build a cover. Let $N_r(S)$ denote the least number of r-balls necessary to open-cover S; the Minkowski-Bouligand dimension of S reduces to:

$$\dim_B(S) = \lim_{r \to 0} \frac{\log(N_r(S))}{\log(1/r)} \tag{7}$$

which clearly reduces to d if S is a d-dimensional manifold embedded in $\mathcal{X}$, so that $N_r(S)$ is roughly proportional to r^{-d} for r small compared to the size of S. From this definition stems the intuition behind dimensionality reduction. Suppose we want to *quantize* $\mathcal{X}$ through a finite number of centroids, such that every point of $\mathcal{X}$ is less than r distant from the closest centroid. This requires $O(r^{-D})$ centroids, which can be overwhelming for real world values of D. Though our actual aim is not to quantise $\mathcal{X}$ but S, whose dimensionality might be significantly lower.

Quantizing through a cover of equal radius balls, as mentioned above, is not necessarily a good idea, in that the ρ-measure of such balls can significantly vary along S. Of course the ρ-measure

of a ball is related to the number of data points that are likely to fall inside that ball, hence, roughly speaking, to the importance of the corresponding centroid.

A *Vector Quantization* procedure is defined by a set of code vectors and an associated partition of the data space [41]. The idea is that compression can be achieved replacing all points in a given cell of a partition by the corresponding code vector. As a matter of fact, Secion 1 describes a vector quantization setting. A variety of different strategies have been proposed, aiming at a quantization that is *even* in some sense, to meet the necessities of the application at hand, either deterministically or stochastically. Some of these will be presented in detail in this thesis, more precisely assuming the error measure of Equation (2) to be convenient.

We review an example that considers a slightly different error measures, but is useful since its optimal solutions are relatively well understood, at least in an asymptotic sense. In this case we talk about *Optimal Quantization*; there are no data involved, partition and centroids are built on ρ to meet the actual minimum of the corresponding objective function. As a result, the only considered partitions are Voronoi-Dirichlet tilings, as optimal quantizers always correspond to such partitions. Indeed, from Equation (2):

$$\mathcal{E}(P_\Lambda) = \sum_{j=1}^{N} \int_{I_j} \|x - c_j\|^2 d\rho(x) \geq \sum_{j=1}^{N} \int_{I_j} \min_l \|x - c_l\|^2 d\rho(x) = \mathbb{E}[\min_l \|x - c_l\|^2], \tag{8}$$

and, by defining $V_l = \{x \in \mathcal{X} | l = \text{argmin}_j \|x - c_j\|\}$, $\Lambda^* = \{V_l\}_l$:

$$\mathcal{E}(P_{\Lambda^*}) = \mathbb{E}[\min_l \|x - c_l\|^2]. \tag{9}$$

On the other hand, at fixed partition $\Lambda^* = \{V_l\}_l$, a minimization with respect to the choice of the centroids factorizes into independent problems, one for each cell:

$$\text{argmin}_{\{c_I \in I\}_{I \in \Lambda^*}} \mathcal{E}(P_{\Lambda^*}) = \left\{ \text{argmin}_{c_I} \mathbb{E}[\|x - c_l\|^2] \right\}_{I \in \Lambda^*},$$

so that, for an optimal quantization, the centroids are always given by the means:

$$c_I = \rho(I)^{-1} \int_I x d\rho(x).$$

Example 2.1.1 ([47]). Let's restrict to the special case of a probability measure ρ that is supported on a compact Riemannian d-submanifold $\mathcal{M} \subset \mathcal{X}$ and admits a density $p : \mathcal{M} \to \mathbb{R}^+$. The following distortion functional is considered:

$$\mathcal{E}_\alpha(P_{S_N}) := \mathbb{E}[\|d_\mathcal{M}(x, P_{S_N}(x))\|^\alpha] \tag{10}$$

with $\alpha > 0$ and $d_\mathcal{M}(,)$ is the Riemannian metric, which is in general larger than the euclidean metric, to an extent that depends on the local curvature. Because [47] is concerned with the asymptotic behaviour of an optimum quantization, they only consider projections P_{S_N} associated with Voronoi-Dirichlet tilings generated by discrete subsets $S_N \subset \mathcal{M}, \#S_N = N$, that is:

$$P_{S_N}(x) = \text{argmin}_{y \in S_N} d_\mathcal{M}(x, y).$$

In this context, for $N \to \infty$, [47] provides an asymptotic characterization of the optimum quantization, together with asymptotics for the corresponding distortion. For clarity we use the notation $\rho^a(A) := \int_A p(x)^a d\omega_\mathcal{M}(x)$ where $d\omega_\mathcal{M}(x)$ is the Riemannian volume measure. Let S_N^* be an optimal quantization, that is, attaining the minimal possible distortion. It can be shown that:

$$\lim_{N \to \infty} \mathcal{E}_\alpha(P_{S_N^*}) N^{-\frac{\alpha}{d}} = C(\alpha, d) \left(\rho^{\frac{d}{\alpha+d}}(\mathcal{M})\right)^{\frac{\alpha+d}{d}} \tag{11}$$

$C(\alpha, d)$ depends only on α and d. Furthermore S_N^* is *asymptotically even* with respect to $\rho^{\frac{d}{\alpha+d}}$ in the following sense:

$$\lim_{N\to\infty} \frac{\#(K \cap S_N^*)}{N} = \frac{\rho^{\frac{d}{\alpha+d}}(K)}{\rho^{\frac{d}{\alpha+d}}(\mathcal{M})} \tag{12}$$

for any measurable[1] $K \subseteq \mathcal{M}$. Roughly speaking, an optimum quantization S_N^* tends to divide $\mathcal{X}$ in cells that have equal mass according to an adjusted measure $\rho^{\frac{d}{\alpha+d}}$.

We will obtain, later in this thesis, results related with (11) from the last example, but in a non-asymptotic framework, that is for finite N, and with a partition that has a constructive definition.

Based upon the insight of this last example, we may step forward from the metric definition of effective dimension in (7), to a measure-based framework.

To this end, let us introduce the Hausdorff dimension. We use a slightly non-standard definition that better suits the context of this section. Equivalence to the usual definition[2] is discussed in [37, page 36]:

$$\mathcal{B}_\delta^\alpha(S) = \inf\left\{ \sum_i r_i^\alpha \mid \{B(x_i, r_i)\}_i \text{ is a cover of } S, r_i < \delta \ \ \forall\ i \right\}, \quad \mathcal{B}^\alpha(S) = \lim_{\alpha\to 0} \mathcal{B}_\delta^\alpha(S). \tag{13}$$

It can be easily shown that $\mathcal{B}^\alpha(S)$ is non-increasing with respect to α, and is always constantly $+\infty$ or 0, except for one value of α in which the discontinuity occurs. This value corresponds to the so called Hausdorff dimension:

$$\dim_H(S) = \sup\{\alpha \mid \mathcal{B}^\alpha = +\infty\} = \inf\{\alpha \mid \mathcal{B}^\alpha = 0\}.$$

If S is a d-dimensional manifold $\dim_H(S) = d$; moreover, we introduced this as a *measure-based* definition because it can be shown that, for $d = \dim_H(S)$, $\mathcal{B}^d(S)$ behaves like a measure[3]. The Hausdorff dimension in fact offers a framework to deal with other measures, in which context a fundamental result is the Mass Distribution Principle; the following is an adaptation from [37, Propositions 4.2,4.9]:

PROPOSITION 2.1.2. *Let $\mathcal{X}$ be a bounded and measurable domain in $\mathbb{R}^D$. Let ρ be a probability measure on $\mathcal{X}$ and let $S \subseteq \mathcal{X}$ be a Borel subset. Suppose that for some $d > 0$ there are $0 < c_1, c_2 < \infty$ and $\epsilon > 0$ such that, $\forall x \in S, r \in (0, \epsilon]$:*

$$c_1 \le \frac{\rho(B(x,r))}{r^d} \le c_2 \tag{14}$$

where $B(x,r)$ is a ball of centrum x and radius r. Then:

$$\frac{\rho(S)}{c_2} \le \mathcal{B}^d(S) \le \frac{2^d}{c_1},$$

and therefore:

$$d = \dim_H(S) \le \dim_B(S).$$

[1] with $\rho(\partial K) = 0$ and compact closure.

[2] as in [37, page 27]

[3] it is indeed proportional to the d-volume measure in the manifold case.

We may ask wether it is really necessary to check the condition (14) on all the open balls of any possible radius and center. Similar situations often occurs in Measure Theory, where large, uncountable families of subsets may sometimes be reduced to countable subfamilies for some specific purpose.

Two facts are clear:

- since the measure is monotone by inclusion, we can restrict to a discrete sequence of radiuses $r_0 > r_1 > \cdots \in (0, R)$ that accumulates on 0;
- it's not necessary to take all possible centers, as long as, for each point $x \in \mathcal{X}$, we have a sequence of nested balls $B(x_1, r_1) \supset B(x_2, r_2) \supset \cdots \ni x$ one for each radius of the sequence.

Notice that, in this way, for each value of r we will have a r-cover of $\mathcal{X}$. As we will see in later sections, these structured families of subsets are tightly related to the so called *dyadic cubes*, and their generalizations, and will be pivotal tools to our main result.

The concept of effective dimensionality is crucial to our analysis, and the assumptions of our main result are tightly related with the Condition (14), with some differences due to the fact that we will be dealing with families of subsets that are not balls nor cubes. These families, called *Partition Trees*, will be introduced in the next chapter. Since we will need Condition (14) to hold true for our main result, it would be useful to derive it from a definition of effective dimension, that is, the opposite implication with respect to Proposition 2.1.2; unfortunately this implication is more delicate, and in fact unavailable at this level of generality. A partial answer is provided by the so called Frostman's Lemma [37]. We will derive a result in this vein, more specific to our approach, in the case in which the support S is a manifold. Condition (14) is in fact equivalent to the so called Ahlfors regularity [46], see Section 5.2.

2.2. Learning the Identity.

Throughout the last subsection we dealt with the problem of quantizing a probability measure ρ, temporarily ignoring the fact that in practice, as anticipated in Section 1, ρ is out of reach and we only have access to a finite sample.

Based of Section 1, since P_Λ is a function, one could be tempted to interpret our problem as inference of a function from a sample, as in Supervised Learning. Indeed Supervised Learning is concerned with the problem of inferring a functional relationship $f : \mathcal{X} \to \mathcal{Y}$ from a dataset of input-output pairs, and based upon minimization of a suitable error measure, such as the least squares loss; see for example [49]. The setting classically consists of a probability measure $\rho_\mathcal{Z}$ on a product space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, with both $\mathcal{X}, \mathcal{Y}$ subsets of finite dimensional Hilbert spaces. A dataset $(\mathbf{x}, \mathbf{y}) = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ of n independent random observations is given, assumed to be identically distributed according to $\rho_\mathcal{Z}$. It is usually assumed that $\mathcal{X}$ is bounded and $\int \|y\|^2 d\rho(x, y) < +\infty$. The marginal probability measure $\rho_\mathcal{X}$ on $\mathcal{X}$ is defined by $\rho_\mathcal{X}(S) := \rho(S \times \mathcal{Y})$. Therefore $d\rho(x, y) = d\rho(y|x) d\rho_\mathcal{X}(x)$, with $\rho(y|x)$ the conditional probability measure on $\mathcal{Y}$ with respect to x.

The goal is estimating the regression function $f_\rho(x)$ defined as the conditional expectation of the random variable y at x:

$$f_\rho(x) := \mathbb{E}(y|x) = \int_\mathcal{Y} y d\rho(y|x).$$

The risk functional:

$$\mathcal{E}_{\mathcal{Z}}(f) := \int_{\mathcal{Z}} |y - f(x)|^2 d\rho,$$

is defined for $f \in L_2(\mathcal{X}, \rho_X)$. This functional has f_ρ as minimizer, since:

$$\mathcal{E}_{\mathcal{Z}}(f) = \mathcal{E}_{\mathcal{Z}}(f_\rho) + \|f - f_\rho\|^2_{L_2(\mathcal{X},\rho_X)}. \tag{15}$$

Analogies between the setting of Section 1 and the one of Equation (15) are not strict, as in the first case there's just one $\mathcal{X}$ space; however, it is in order to notice that the least squares distortion in Equation (2) is such that the identity function $I : \mathcal{X} \to \mathcal{X}, x \mapsto x$ gives clearly zero distortion. This suggests that we are in a peculiar supervised learning scenario in which we know the *target function* in advance.

Based upon this insight, we can define from (15) another risk functional:

$$\mathcal{E}_{\mathcal{X}}(f) := \|f - f_\rho\|^2_{L_2(\mathcal{X},\rho_X)} = \mathcal{E}_{\mathcal{Z}}(f) - \mathcal{E}_{\mathcal{Z}}(f_\rho)$$

which is still minimized by f_ρ, but doesn't explicitly mention the space $\mathcal{Y}$.

Hence our objective is still to find an estimator $\widehat{f}$ for f_ρ based on the data $\{x_1, \dots, x_n\}$ such that the quantity $\mathcal{E}_{\mathcal{X}}(\widehat{f})$ is small with high probability. Because a solution is known, the problem might seem trivial, but in fact its formulation is not complete yet.

Indeed, as in Supervised Learning, we need to sets a suitable hypothesis space $\mathcal{H}$, in such a way that the solution $\widehat{f} \in \mathcal{H}$ whilst having low expected risk, also exhibits properties that make it convenient in practice in some sense. Since f_ρ is known it is not informative, and as a result it is natural to choose $\mathcal{H}$ such that $f_\rho \notin \mathcal{H}$; therefore the risk functional cannot vanish on $\mathcal{H}$ and its analysis is non trivial.

We already outlined in Section 1 what kind of hypothesis space we are looking for. Indeed the P_Λ functions are piece-wise constant on certain partitions Λ with a limited number of cells/centroids. This choice is convenient as it provides useful information about the dataset as a whole, in a *parsimonious* way. Given a partition of $\Lambda = \{I_j\}_{j=1,\dots,N}$, the space:

$$\mathcal{H}_\Lambda = \{f : \mathcal{X} \to \mathcal{X} \mid f(x) = \sum_{j=1}^{N} c_j \mathbb{1}_{I_j}(x), c_j \in \mathbb{R}^D \ \forall j\}$$

is a vectorial space of dimension ND, even though, to any practical use, only the functions satisfying $c_j \in I_j, \forall j$ are meaningful. Furthermore the partition will not be fixed, so we deal with a function space $\mathcal{H}$ that is not a vector space, but rather a union of patches, each of which is a subset of a vector space corresponding to an admissible partition Λ. Admissible partitions will have constraints, whose operative definition is provided in later sections.

There is a huge literature about parsimonious representation of functions through multi-resolution analysis and related topics. Our work can be regarded as an attempt to import some of this knowledge in our different setting, as described in the next section.

2.3. Multi-Resolution and Adaptivity.

Wavelets, in their numerous instances and developments, are ubiquitous in signal processing, see for example [28, 64]. Their strength stems from the hierarchical organization of the information

contained in a signal, so that simple elementary components can be added up to a desired level of accuracy in a way that is efficient and manageable; furthermore they offer an intuitive tool for alternative definitions of important function spaces, such as L_p, H_p and Besov spaces. We provide here a very brief example that will turn out useful as reference to justify some later choices.

Consider the unit interval $(0,1] \subset \mathbb{R}$ with the Lebesgue measure. Within the unit interval one can build the family of subsets:

$$\mathbb{T} = \{I_{j,k} = 2^{-j}(k + (0,1])\}_{k\in\{1,\dots,2^j\}, j\in\{0,1,\dots\}}$$

also called dyadic intervals. Clearly, for each value of j, the subset $\{I_{j,k}\}_{k\in\{1,\dots,2^j\}}$ is a partition of the unit interval. It is well known that, provided appropriate orthogonal functions $\phi, \psi \in L_2((0,1])$ supported on $(0,1]$, the set of functions:

$$\{\phi(\cdot), 2^{j/2}\psi(2^j \cdot -k)\ \forall k \in \{1,\dots,2^j\}, j \in \{0,1,...\}\} \tag{16}$$

is an orthonormal basis for $L_2((0,1])$, with each ψ_I supported on a dyadic interval $I \in \mathbb{T}$. An example is provided by the *Haar wavelets*: $\phi_{(0,1]} = \mathbb{1}_{(0,1]}, \psi_{(0,1]} = \mathbb{1}_{(0,1/2]} - \mathbb{1}_{(1/2,1]}$.

For expository purposes we restrict our attention to orthonormal wavelets. In this case, any function $f \in L_2((0,1])$ will be represented by a set of *wavelet coefficients*:

$$\begin{aligned} a_0(f) = \langle \phi, f\rangle, \{a_I(f) = \langle \psi_I, f\rangle\}_{I\in\mathbb{T}} \\ f = a_0(f)\phi + \sum_{I\in\mathbb{T}} a_I(f)\psi_I \end{aligned} \tag{17}$$

whose decay with respect to j carries important information about the smoothness of f; furthermore one can consider subspaces of $L_2((0,1])$ given by:

$$V_n = \text{span}\left(\phi, \{\psi_{j,k}\}_{k\in\{1,\dots,2^j\}, j\in\{0,\dots,n\}}\right),$$

being linear finite-dimensional subspaces, so that $P_{V_n}(f)$, the L_2-orthogonal projection of f onto V_n, provides an approximation of f at resolution $j = n$. Linearity clearly entails:

$$\inf_{g\in V_n} \|f - g\|^2_{L_2(0,1]} = \|f - P_{V_n}(f)\|^2_{L_2(0,1]} = \sum_{j=n+1}^{\infty}\sum_{k=0}^{2^j-1} |a_{j,k}(f)|^2.$$

Oftentimes functions are not evenly smooth on their domain, making V_n approximations inconvenient; this motivated the introduction of non-linear approximation strategies. One of these consists in considering nonlinear subspaces of $L_2((0,1])$ given by:

$$\Sigma_\eta = \{f \in L_2((0,1])|\forall I \in \mathbb{T}, a_I(f) \notin (0,\eta)\}$$

Being Σ_η a non-linear space, the L_2-best approximation of a function f in Σ_η cannot be found by orthogonal projection; still this can be found by simply keeping the only non-zero wavelet components such that $a_I(f) \geq \eta$, that is, by *wavelet thresholding*; indeed, by orthonormality:

$$\inf_{g\in\Sigma_\eta} \|f - g\|^2_{L_2(0,1]} = \sum_{a_I(f)<\eta} |a_I(f)|^2.$$

Given an approximation $g \in \Sigma_\eta$ of $f \in L_2((0,1])$, the subset $S_\eta(g) \subset \mathbb{T}$ whose $I \in S_\eta(g)$ yield $a_I(g) \neq 0$ can be arbitrarily unstructured, and this, as observed in [24], can cause inconvenient storage of the coefficients. Therein, it is proposed to choose a representation of f that is slightly less compact, in terms of number of wavelet components, but with constraints on the structure of the subset $\mathcal{T}_\eta \subset \mathbb{T}$ whose wavelet coefficients are preserved.

Consider $\mathbb{T}$ as a graph, in fact a tree, whose root is $I_{0,0}$, and each edge connects an interval $I_{j,k}$ with a subinterval of the form $I_{j+1,k'}$. The constraint on $\mathcal{T}_\eta$ requires it to be a subtree of $\mathbb{T}$ containing the root, and actually the smallest such subtree containing S_η. A similar strategy is adopted in [12] as well, for different reasons, and is also at the core of our work. This point is further discussed in Section (3.1).

The idea of a multiscale approach to vector quantization is not new [41, page 410], and tree based algorithms are a key tool to reduce complexity of countless data mining tasks, such as nearest neighboor search. Yet theoretical estimates of the L^2-distortion, in the form of learning rates, were not available in this case, to the best of our knowledge.

A set of functions like the wavelet basis (16) is a prototypical example of what is more generally called a *Dictionary*. Any function can be represented on any basis, but the corresponding coefficients will be in general all non-zero. A dictionary is a (typically overcomplete) set of functions, designed in such a way that a specific set of functions of interest can be represented with *sparse* coefficients, that is, only using a small subset of the dictionary. Sparse Dictionary Learning tries to find a dictionary that allows the sparsest representations for a given set of functions; see for example [56]. In this case the representation is called *adaptive*, as it adapts to a specific source. Vector quantization, as in Equation (1), can be considered an extreme case of sparse dictionary learning, in which all coefficients are zero but one; see for example [77]. We will touch again upon the subject through next subsection.

2.4. Higher Order.

In Section 1 we outlined the general form that the projection P_Λ will take for the main algorithm we analyze, as a piece-wise constant function, that is, whose image is a finite set of centroids, one for each cell of the partition Λ. We mentioned that this point of view is shared with Vector Quantization algorithms [41], being k-means perhaps the most popular instance of such [5].

In Approximation Theory from a piece-wise constant approximation we can step further to piece-wise linear or piece-wise polynomial approximations. In the same way, here we might consider functions P_Λ that project each cell of the partition Λ onto a corresponding hyperplane, rather than to a centroid; that would be considered a piece-wise linear approximation.

Let $c_1, \dots, c_N$ be a set of centroids as before, and $V_1, \dots, V_N$ be a set of d dimensional linear subspaces, where d is assumed fixed and known, and $\Lambda = \{I_1, \dots, I_N\}$ defines a partition of $\mathcal{X}$ as before. Then, the nonlinear projection will, in this case, take the form

$$P^1_\Lambda(x) = \sum_{j=1}^{K} \mathbb{1}_{I_j}(x)\big(c_j + P_{V_j}(x - c_j)\big), \tag{18}$$

therefore orthogonally projecting each cell I_j onto the corresponding affine hyperplane, parallel to V_j and through c_j. A natural choice for empirically estimating c_j and V_j is via mean and PCA, both restricted on the cell labeled by j. In the aforementioned manifold case, the hyperplane parallel to V_j through c_j can be thought of as a rough approximation of a tangent space to the manifold at some point on it, within I_j. However c_j is in general not on the manifold, and the hyperplane in (18) is not meant to approximate a tangent space but the manifold itself (its portion lying inside I_j). The two hyperplanes may not even be parallel, in particular when the curvature is not isotropic.

While Equation (1) corresponds to the Vector Quantization problem, here Equation (18) might be considered a special instance of the Dictionary Learning or of the Subspace Clustering problem. Some popular algorithms addressing the latter problems are k-SVD [2] and k-flats [19], even though none of them exactly uses the form (18). On the other hand a projection of the form (18) is in fact at the base of the algorithm GMRA introduced in [3, 63, 59]. The latter also shares with our approach the multi-scale framework, that was suggested through the previous sections and will be further clarified later. For this reason, our approach might be considered a 0-th order version of the GMRA, but comparison should be carried out carefully. To this respect we draw the reason's attention on two points:

- Working with a representation of the form (18) requires stronger prior knowledge about the probability measure ρ. Algorithmically, the intrinsic dimension d is either assumed to be known, or has to be inferred, while (1) does not require that. From the point of view of the statistical guaranties, GMRA analysis also requires more algorithm-specific assumptions with respect to our approach, as we will comment in detail later.
- By comparing the expected distortion caused by (1) or (18), it's only meaningful to compare representations that carry the same amount of information, and that have similar computational complexity. For instance, performing (18) requires an amount of extra storing space that is proportional to dDN, if the two representations work with the same cardinality N. On the other hand, beside the distortion, a more expressive representation might be useful to some specific purpose, for example for representing of functions over $\mathcal{X}$; we will not cover this subject any further.

Also the two algorithms build the partition Λ in similar yet different ways. Detailed comparison of algorithms and statistical guaranties is presented in Section 4.

The empirical evaluation of the intrinsic dimension is a topic worth interest in its own right. In fact the GMRA algorithm originally stemmed from the observation that the procedure used to empirically estimate the hyperplanes V_j in (18), the Principal Component Analysis (PCA) [13], naturally offers a mean towards dimensionality estimation. PCA works on the *empirical covariance matrix* of a point cloud, in our case the portion of the dataset that lies inside a cell. The covariance matrix of a distribution is known to have eigenvalues that gauge how far the distribution spreads along each of the principal (orthogonal) axes (variance-wise). Given that a cell is comparable to a ball of radius r, it can be shown that the d leading eigenvalues are expected to be proportional to r, while the remaining tend to vanish as $O(r^2)$, so that, at an appropriate scale r, a gap opens in the spectrum, suggesting the effective dimension d. We refer to [60] and reference therein for pointers to the specific literature.

2.5. Manifold Ansatz.

In this section some loosely related literature about data representation will be presented. This was mostly motivational to GMRA [3, 63, 59] and therefore for us.

We already mentioned multiple times, as a special case, the possibility that the support of ρ may happen to be a smooth manifold embedded in $\mathcal{X}$; this is indeed a case in which several computations can be carried out neatly; see Proposition 3.4.1. On the other hand this case also turns out to be particularly relevant in several applications. In the literature the attention is sometimes devoted to finding an approximation for the support of the distribution ρ rather

than for the distribution itself; hence we can talk about Set Learning or, more often, Manifold Learning; referring to data coming from various applications, it has been noticed that, except for the possible presence of some sort of noise, the support of the data seems to be often well-approximated by a smooth low-dimensional manifold, embedded in the ambient space. The latter idea goes under the name of Manifold Hypothesis; roughly speaking, naturally occurring data are most likely generated by hidden dynamical systems that have much fewer degrees of freedom than the actual dimension of the data.

Some effort has been devoted to formalize when the manifold hypothesis is quantitatively justified; for example [38] and [73] proved that, depending on a desired confidence and under suitable assumptions, there is critical value of the sample size over which it is possible to say whether the data describe a manifold in a specific geometrical sense, while [76] derived the critical sample sizes over which we can infer topological invariants of the manifold underlying the data with high confidence; remarkably, these results are independent from the context the data come from. Arguably the research about convergence of the Graph Laplacian of a point cloud to the Laplace-Beltrami operator of the underlying manifold [50] also belongs in this vein.

Once the Manifold Hypothesis is assumed holding true (or taken as a convenient ansatz) an algorithm is then needed to construct an empirical estimator of the supposed manifold support of the data.

Elementary empirical estimators are of course the mean and the PCA, used when the model consists of a flat manifold, that is, an affine subspace (a single point being considered a 0-dimensional instance of such). Statistical properties of mean and PCA are well understood [13], providing finite-sample concentration inequalities.

Starting from PCA a number of developments can be considered to relax the linearity assumption; these include Isomap [102], Hessian Eigenmaps [32], Laplacian Eigenmaps [7] and related developments such as Diffusion Maps [70]. Another way to tackle non-linearity is kernel PCA, performing PCA after a suitable nonlinear embedding [91, 90]. We also mention Principal Manifolds [44], the Autoencoders designed for non-linear dimensionality reduction [9], Local Linear Embedding [88] and Local Tangent Spaces Alignment [111], for all of which we refer the reader to the specific literature.

Statistical guaranties about the performance of the algorithms above are rarely available.

Vector Quantization in general, as well as the piece-wise linear approach mentioned in the previous subsection, can be considered instances of Manifold Learning. Indeed this is the context in which GMRA was convieved.

Some of the algorithms above are considered *global*, especially kernel-PCA and the ones involving a Laplacian operator; indeed they are meant to process the dataset in its entirety, aiming at a representation that takes into account both large scale features and details. This typically implies at some point the necessity to invert matrices as large as the sample size, which is considered computationally intensive for real world data. On the other hand a method would be considered *local* if it splits the dataset in smaller batches to be processed independently. From the methods above some can be considered mixed local, as they entail both local and global steps, cooperating to balance manageability and accuracy. That's in fact the case for the prototypes given in Equation (1) and (18); the representation itself is local in structure, except that we haven't mentioned yet how the partitions Λ are sorted out.

3. Algorithm Statement and Main Result

In Section 1 we outlined the problem of our interest, as constructing en empirical estimator for a Vector Quantization, using multi-resolution tools. A set of centroids and a corresponding partition are considered, and then all data points in each cell of the partition are represented by the corresponding centroid. In this section we provide the operative definition of an algorithm addressing this problem, and it's statistical analysis.

The most classical approach in this context is probably K-means, where the set centroids (means) is defined by a non-convex optimization over all possible partitions of the data; the latter has been extensively studied and found countless applications, a comparison is presented in Section 4, together with review of some other related subjects. Our approach is not to be considered a mere alternative to K-means, as it offers a multi-resolution representation of the dataset; the hierarchical organization of the information naturally provides more flexibility for exploratory purposes.

The novelty of the present work stems from framing the vector quintization problem within a supervised learning setting, by providing a multi-resolution empirical estimator, and therein obtaining statistical guaranties in the form of learning rates, stated from simpler assumptions as compared to similar works in the literature; furthermore non-asymptotic estimates for the expected distortion of the corresponding infinite sample problem are provided.

3.1. Partition Trees and Subtrees.

In Section 1 a one-parameter family of partitions Λ_η was introduced, such that Λ_{η_2} is a refinement of Λ_{η_1} whenever $\eta_2 < \eta_1$. An instance of such a family of partitions was not provided yet, except for the dyadic cubes used to define wavelets. Before describing the family Λ_η that will be the subject of our main result, we point out here that Λ_η may or may not depend on ρ. Of course a depending one has higher chances to provide higher accuracy at lower complexity (that is, at lower N). On the other hand a Λ_η depending on ρ, i.e. *adaptive*, will in turn be more expensive to build.

The strategy that we pursue, borrowed from [12] and references therein, is to start with a family of partitions that are easy to build, independent or loosely dependent[4] on ρ, that will be called *partition tree* $\mathbb{T}$. Then from this partition tree, an adaptive family Λ_η will be extracted, with an algorithm that depends on ρ, and aims at produce a low expected distortion; in fact it will depend only on the dataset $x_1, \dots, x_n$, as ρ would in practice be unknown.

In Section 2.3 we anticipated as a constrained subset of the dyadic cubes can provide an efficient representation in the context of wavelet thresholding; the constraint consists in the subset being a subtree. We will show that a similar constraint on subsets of a partition tree $\mathbb{T}$ in fact naturally provides a partition, and therefore is particularly convenient to tackle partition refinements in our context. As anticipated, also the way the subtree is selected will be tightly related with the wavelet thresholding introduced in Section 2.3.

[4] a loose dependence is implicit in our Assumption 3.3.1

We begin introducing the definition of a partition tree. In the following we denote by $\mathcal{X} \subset \mathbb{R}^D$ the data space endowed with its natural Borel σ-algebra $\mathcal{B}(\mathcal{X})$ and by $\sharp A$ the cardinality of a set A.

DEFINITION 3.1.1. *A partition tree $\mathbb{T}$ is a denumerable family $\{\Lambda_j\}_{j\in\mathbb{N}}$ of partitions of $\mathcal{X}$ satisfying*

a) $\Lambda_0 = \{\mathcal{X}\}$ is the root of the tree;
b) each family $\Lambda_j = \{I\}_{I\in\Lambda_j}$ is a finite partition of $\mathcal{X}$ of Borel subsets, i.e

$$\begin{cases} \mathcal{X} = \bigcup_{I\in\Lambda_j} I \\ I \cap J = \varnothing \qquad \forall I, J \in \Lambda_j, I \neq J \\ \sharp\Lambda_j < +\infty \\ I \in \mathcal{B}(\mathcal{X}) \qquad I \in \Lambda_j \end{cases} ;$$

c) for each $I \in \Lambda_j$, there exists a family $\mathcal{C}(I) \subseteq \Lambda_{j+1}$ such that

$$\begin{cases} I = \bigcup\limits_{J\in\mathcal{C}(I)} J \\ \sharp\mathcal{C}(I) \leq a \end{cases} , \tag{19}$$

where $a \in (0, +\infty)$ is a constant depending only on $\mathbb{T}$.

The family $\mathbb{T}$ is called tree because it is naturally associated with a graph, in fact a tree. This graph has a vertex corresponding to each cell of all partitions, and edges corresponding to couples of cells $I, J \in \mathbb{T}$ such that $I \supseteq J$ and such that the corresponding partitions are consecutive in the refinement sequence. This will be alternatively denoted as a family Λ_j of partitions, in which j is called *depth*, that is the edge-counting distance from Λ_0, the coarsest partition made by only one cell.

Note that, we allow the partition tree to have arbitrary, possibly infinite, depth, needed to derive asymptotic results.

Further, notice that, since $\#\Lambda_j \leq a^j$ the constant a characterize how the cardinality of each partition increases at finer scale. The case $a = 2$ corresponds to dyadic trees.

We add some further definitions. For any $j \in \mathbb{N}$, and $I \in \Lambda_j$ the depth of I is j and is denoted by j_I. The cells in $\mathcal{C}(I) \subset \Lambda_{j+1}$ are the children of I, the unique cell $J \in \Lambda_{j-1}$ such that $I \in \mathcal{C}(J)$ is the parent of I and is denoted by $\mathcal{P}(I)$ (by definition $\mathcal{P}(\mathcal{X}) = \mathcal{X}$). We regard $\mathbb{T}$ as a set of nodes where each node is defined by a cell I with its parent $\mathcal{P}(I)$ and its children $\mathcal{C}(I)$. The following definition will be crucial.

DEFINITION 3.1.2. *A (proper) subtree of $\mathbb{T}$ is a family $\mathcal{T} \subset \mathbb{T}$ of cells such that $\mathcal{P}(I) \in \mathcal{T}$ for all $I \in \mathcal{T}$ and*

$$\Lambda(\mathcal{T}) = \{I \in \mathbb{T} : I \notin \mathcal{T}, \mathcal{P}(I) \in \mathcal{T}\},$$

denotes the set of outer leaves.

It is important in what follows that $\Lambda(\mathcal{T})$ is a partition of $\mathcal{X}$ if $\mathcal{T}$ is finite, see Lemma 5.1.2. If the $\mathcal{T}$ is not finite $\Lambda(\mathcal{T})$ can still define a partition of unbounded cardinality (enumerable), but this case is beyond the scope of our present work.

3.2. Reconstruction Trees.

We next discuss a data driven procedure to derive a suitable partition and a corresponding nonlinear projection. To do this end, we need a few definitions depending on an available dataset $x_1, \dots, x_n$.

For each cell I, we fix an arbitrary point $\hat{x}_I^* \in \mathcal{X}$ and define the corresponding cardinality and center of mass, respectively, as

$$n_I = \sum_{i=1}^{n} \mathbb{1}_I(x_i), \qquad \widehat{c}_I = \begin{cases} \dfrac{1}{n_I} \displaystyle\sum_{i=1}^{n} x_i \mathbb{1}_I(x_i) & \text{if } x \in I \text{ and } n_I \neq 0 \\ \hat{x}_I^* & \text{if } x \in I \text{ and } n_I = 0 \end{cases},$$

where $\mathbb{1}_I$ is the characteristic function of I, *i.e.*

$$\mathbb{1}_I(x) = \begin{cases} 1 & x \in I \\ 0 & x \notin I \end{cases}.$$

If $0 \in \mathcal{X}$, a typical choice is $\hat{x}_I^* = 0$ for all cells $I \in \mathbb{T}$. While $\mathcal{E}(\hat{P}_n)$ depends on the choice of $\hat{x}_I^*$, our bounds hold true for all choices. We point out that it is more convenient to choose $\hat{x}_I^* \in I$, as this (arbitrary) choice produces an improvement of $\mathcal{E}(\hat{P}_n)$ for free, in particular whenever $\rho_I > 0$ but $n_I = 0$.

Using this quantity we can define a local error measure for each cell I,

$$\widehat{\mathcal{E}}_I = \frac{1}{n} \sum_{x_i \in I} \|x_i - \widehat{c}_I\|^2 = \frac{1}{n} \sum_{i=1}^{n} \|x_i - \widehat{c}_I\|^2 \, \mathbb{1}_I(x_i)$$

as well as the potential error difference induced by considering a refinement,

$$\widehat{\epsilon}_I^2 = \widehat{\mathcal{E}}_I - \sum_{J \in \mathcal{C}(I)} \widehat{\mathcal{E}}_J = \frac{1}{n} \sum_{J \in \mathcal{C}(I)} \|\widehat{c}_J - \widehat{c}_I\|^2, \tag{20}$$

where the second equality is consequence of the between-within decomposition of the variance. A motivation for considering a thresholding on $\widehat{\epsilon}_I^2$ has been mentioned in Section (2.3). Indeed, referring to the notation therein, considering the special case of Haar wavelets, we have:

$$\phi(\cdot) = \mathbb{1}_{(0,1]}(\cdot), \psi(\cdot) = \phi(\cdot) - \phi(2\cdot).$$

For any interval $I_{j-1,k} =: I$ and its subintervals $I_{j,2k-1} =: J_1, I_{j,2k} =: J_2$, it easily follows:

$$|a_{J_1}|^2(f) + |a_{J_2}|^2(f) = \|f - f_I\|^2_{L^2(0,1]} - \left(\|f - f_{J_1}\|^2_{L^2(0,1]} + \|f - f_{J_2}\|^2_{L^2(0,1]} \right) := \epsilon_I^2(f)$$

where $f_I = \int_I \frac{f(x)}{|I|} dx$, and $|I|$ is the Lebesgue measure of I. Our algorithm is based on thresholding a quantity formally corresponding to $\epsilon_I(f)$, which in turns is very close to thresholding the actual wavelet coefficients, with the difference required by the necessity of building a partition of $\mathcal{X}$.

Following [12], we first truncate the partition tree at a given depth, depending on the size of the data set. More precisely, given $\gamma > 0$, we set

$$j_n = \left\lfloor \frac{\gamma \ln n}{\ln a} \right\rfloor \qquad \Longrightarrow \qquad a^{j_n} \leq n^{\gamma}. \tag{21}$$

We add a a few comments, the above vector quantization procedure, that we call reconstruction tree, is recursive and depends on the threshold η. Different quantizations and corresponding distortions are achieved by different choices of η. Smaller values of η correspond to vector quantization quantizations with smaller distortion. It is a clear that the empirical distortion becomes zero for a suitably small η corresponding to having a single point in each cell. Understanding the behaviour of the expected distortion as function of η and the number of points is our main theoretical contribution. Before discussing these results we discuss the connection of the above approach to related ideas. A similar construction is given in [59], where however the thresholding criterion η depends on the scale, see Section 2.3 of the cited reference.

The parameter γ is an a-priori bound on the depth of the partion tree. It is needed to show that $P_{\mathcal{T}_\eta}$ is closed to its empirical counterpart $\widehat{P}_{\mathcal{T}_\eta}$ and it controls the size of the truncated partition tree $\mathcal{T}_n^*$, see Proposition 7.2.4 and bound (82). For essentially the same technical reason, in [59] the partition tree is truncated to the largest subtree whose leaves contain at least d points and d denotes the intrinsic dimension of the data. From a pratical point of view, the parameter γ allows to have partition $\widehat{\Lambda}_n$ with a small number of cells, so that the final vector quantization is parsimonious.

Deeper trees are considered as data size grows. This choice is only meant to prevent our statistical analysis to be affected by unlikely worst case scenarios, and introduces no substantial restriction from a practical point of view.

As a second step, we select the cells such that $\hat{\epsilon}_I \geq \eta$. Since $\hat{\epsilon}_I$ is not an decreasing function with the depth of the tree, this requires some care – see Remark 5.1.1 for an alternative construction. Indeed, for a threshold $\eta > 0$, we define the subtree

$$\widehat{\mathcal{T}}_\eta = \begin{cases} \{\mathcal{X}\} & \text{if } \widehat{\epsilon}_I < \eta \quad \forall I \in \bigcup_{j \leq j_n} \Lambda_j \\ \{I \in \mathbb{T} \mid \exists j \leq j_n,\, J \in \Lambda_j,\, J \subset I,\, \widehat{\epsilon}_J \geq \eta\} & \text{otherwise} \end{cases} \tag{22}$$

and $\widehat{\Lambda}_\eta$ is defined as outerleaves of $\widehat{\mathcal{T}}_\eta$, *i.e.* $\widehat{\Lambda}_\eta = \Lambda(\widehat{\mathcal{T}}_\eta)$, see Figure 1 below. Note that $\widehat{\mathcal{T}}_\eta$ is finite, so that by Lemma 5.1.2 $\widehat{\mathcal{T}}_\eta$ is a partition of $\mathcal{X}$ such that $j_I \leq (\gamma \ln n)/\ln a$ for all $I \in \widehat{\Lambda}_\eta$.

The code vectors are the centers of mass of the cells the above empirical partition, and the corresponding nonlinear projection is

$$\widehat{P}_\eta = \sum_{I \in \widehat{\Lambda}_\eta} \hat{c}_I \mathbb{1}_I(x). \tag{23}$$

3.3. General Assumptions.

In this section, we introduce our main assumptions and then discuss a motivating example where data are sampled at random from a manifold.

We consider a statistical learning framework, in the sense that we assume the data to be random samples from an underlying probability measure. More precisely, we assume the available data to be a realization of n identical and independent random vectors $X_1, \ldots, X_n$ taking values in a bounded subset $\mathcal{X} \subset \mathbb{R}^D$ and we denote by ρ the common law. Up to a rescaling and a translation, we assume that $0 \in \mathcal{X}$ and

$$\operatorname{diam}(\mathcal{X}) = \sup_{x,y \in \mathcal{X}} \|x - y\| \leq 1. \tag{24}$$

Our main assumption relates the distribution underlying the data to the partition tree to be used to derive a MSVQ via reconstruction trees. To state it, we recall the notion of essential diamater of a cell I, namely

$$\mathrm{diam}_\rho(I) = \inf_{\substack{J \subset I \\ \rho(J)=0}} \mathrm{diam}(I \setminus J).$$

ASSUMPTION 3.3.1. *There exists $s > 0$ and $b > 1$ such that for all $I \in \mathbb{T}$*

$$\mathrm{diam}_\rho(I) \leq C_1 \rho(I)^s \tag{25a}$$

$$\mathrm{diam}_\rho(I) \leq C_2 b^{-j_I} \tag{25b}$$

where $C_1 > 0$ and $C_2 > 0$ are fixed constants depending only on $\mathbb{T}$.

To simplify the notation, we write $c_{\mathbb{T}}$ for a constant depending only on s, b, C_1, C_2 and we write $A \lesssim B$ if there exists a constant $c_{\mathbb{T}} > 0$ such that $A \leq c_{\mathbb{T}} B$.

Given the partition tree $\mathbb{T}$, the parameters s and b define a class $\mathcal{P}_{b,s}(\mathbb{T})$ of probability measures ρ and for this class we are able to provide a finite sample bound on the distortion error of our estimator $\hat{P}_\eta$, see (29). In the context of supervised machine learning $\mathcal{P}_{b,s}(\mathbb{T})$ is an a-priori class of distributions defining a upper learning rate, see (31a). It remains an open problem to provide a lower min-max learning rate.

Clearly, (25b) is implied by the distribution-independent assumption that

$$\mathrm{diam}(I) \lesssim b^{-j_I} \qquad \text{for all } I \in \mathbb{T}, \tag{26}$$

i.e. the diameter of the cells goes to zero exponentially with their depth. This assumption ensures that the reconstruction error goes to zero and, in supervised learning, it corresponds to the assumption that the hypotheses space is rich enough to approximate any regression function, compare with condition (A4) in [59].

Eq. (25a) is a sort of regularity condition on the shape of the cells and, if it holds true, (25b) is implied by the following condition

$$\rho(I) \lesssim c^{-j_I} \qquad \text{for all } I \in \mathbb{T}, \tag{27}$$

which states that the volume of the cells goes to zero exponentially with their depth.

3.4. Results and Constructions about Partition Trees.

The parameter s is related to the intrinsic dimension of the data, still depending also on the partition tree. For example, if $\mathcal{X} = [0,1)^D$ is the unit cube and ρ is given by

$$\rho(E) = \int_E p(x)dx \qquad E \in \mathcal{B}(\mathcal{X}),$$

where dx is the Lebesgue measure of $\mathbb{R}^D$ and the density p is bounded from above and away from zero, see (28b) below, it is easily to check that the family $\mathbb{T} = \{\Lambda_j\}$ of dyadic cubes

$$\Lambda_j = \{[2^{-j}(k_1 - 1), 2^{-j}k_1) \times \ldots \times [2^{-j}(k_D - 1), 2^{-j}k_D) \mid k_1, \ldots, k_D = 1, ..., \ldots, 2^j\} \qquad j \in \mathbb{N}$$

is a partition tree satisfying Assumption 3.3.1 with $s = 1/D$ and a suitable $b > 1$. The construction of dyadic cubes can be extended to more general settings, see [23, 42] and references

therein, by providing a large class of other examples, as shown by the following result. The proof is deferred to Section 5.

PROPOSITION 3.4.1. *Assume that the support $\mathcal{M}$ of ρ is a connected submanifold of $\mathbb{R}^D$ and the distribution ρ is given by*

$$\rho(E) = \int_{E\cap\mathcal{M}} p(x) d\rho_{\mathcal{M}}(x) \qquad E \in \mathcal{B}(\mathcal{X}) \tag{28a}$$

$$0 < p_1 \leq p(x) \leq p_2 < +\infty \qquad x \in \mathcal{M} \tag{28b}$$

where $\rho_{\mathcal{M}}$ is the Riemannian volume element of $\mathcal{M}$, then there exists a partition tree $\mathbb{T}$ of $\mathcal{X}$ satisfying Assumption 3.3.1 with $s = 1/d$, where d is the intrinsic dimension of $\mathcal{M}$.

We recall that, as a submanifold of $\mathbb{R}^D$, $\mathcal{M}$ becomes a compact Riemannian manifold with Riemannian distance $d_{\mathcal{M}}$ and Riemannian volume element $\rho_{\mathcal{M}}$, see for example [82]. We stress that the construction of the dyadic cubes only depend on $d_{\mathcal{M}}$.

By inspecting the proof of the above result, it is possible to show that a partition tree satisfying Assumptions 3.3.1 always exists if there are a metric d and a Borel measure ν on $\mathcal{M}$ such that $(\mathcal{M}, d, \nu)$ is an Ahlfors regular metric measure [46, page 413], ρ has density p with respect to ν satisfying (28b) and the embedding of $(\mathcal{M}, d)$ into $(\mathbb{R}^d, \|\cdot\|)$ is a Lipschitz. function.

This proposition is only a step towards understanding the geometrical meaning of our assumptions and their algorithmic implications. In Section 4 We provide further discussion, together with some options for the practical construction of a partition tree satisfying the requirements.

3.5. Main Result.

In this section we state and discuss our main results, characterizing the expected distortion of reconstruction trees, as well as the result on the Approximation Error. The proofs are deferred to Sections (5)(7). Our first result is a probabilistic bound for any given threshold η. Recall that $s > 0$ is defined by (25a) and $\hat{P}_\eta$ by (22) and (23).

THEOREM 3.5.1. *Fix $\gamma > 0$ as in (21) and $\eta > 0$, for any $0 < \sigma < s$*

$$\mathbb{P}\left[\mathcal{E}[\widehat{P}_\eta] \gtrsim \eta^{\frac{4\sigma}{2\sigma+1}}(1+t)\right] \lesssim \eta^{-\frac{2}{2\sigma+1}} \exp\left(-c_{\mathbb{T}} n\eta^2 t\right) + (n^\gamma + \eta^{-\frac{2}{2\sigma+1}}) \exp\left(-c_a n\eta^2\right) \qquad t > 0, \tag{29}$$

where $c_a = \frac{1}{128(a+1)}$ and $c_{\mathbb{T}} > 0$ depends on the partion tree $\mathbb{T}$.

As shown in Remark 5.1.10, it is possible to set $\sigma = s$ up to an extra logarithmic factor.

Next, we show how it allows derive the best choice for η as a function of the number of examples, and a corresponding expected distortion bound.

COROLLARY 3.5.2. *Fix $\gamma > 1$, $\beta > 0$ and set*

$$\eta_n = \sqrt{\frac{(\gamma+\beta)\ln n}{c_a n}} \qquad \text{and} \qquad \widehat{P}_n = \widehat{P}_{\eta_n} \qquad n \geq 1, \tag{30}$$

where $c_a = \frac{1}{128(a+1)}$. *Then for any* $0 < \sigma < s$

$$\mathbb{P}\left[\mathcal{E}[\widehat{P}_n] \gtrsim \left(\frac{\ln n}{n}\right)^{\frac{2\sigma}{2\sigma+1}} (1+t)\right] \lesssim \frac{1}{n^{\beta}} + \frac{1}{n^{\bar{c}_{\mathbb{T}} t - 1}}, \tag{31a}$$

where $\bar{c}_{\mathbb{T}} > 0$ *is a constant depending on the partition tree* $\mathbb{T}$. *Furthermore*

$$\lim_{t\to+\infty} \limsup_{n\to+\infty} \sup_{\rho\in\mathcal{P}_{b,s}(\mathbb{T})} \mathbb{P}\left[\mathcal{E}[\widehat{P}_n] \gtrsim \left(\frac{\ln n}{n}\right)^{\frac{2\sigma}{2\sigma+1}} t\right] = 0, \tag{31b}$$

where $\mathcal{P}_{b,s}(\mathbb{T})$ *is the family of distributions* ρ *such that Assumptions 3.3.1 hold true.*

If t is chosen large enough so that $\bar{c}_{\mathbb{T}} t - 1 = \beta$, then bound (31a) reads as

$$\mathbb{P}\left[\mathcal{E}[\widehat{P}_n] \geq c_1 \left(\frac{\ln n}{n}\right)^{\frac{2\sigma}{2\sigma+1}}\right] \leq c_2 \frac{1}{n^{\beta}}$$

where c_1 and c_2 are suitable constants depending on $\mathbb{T}$. Comparison of this result with other similar in the literature is presented in Section 4.

The proof of Theorem 3.5.1 relies on splitting the error in several terms. In particular, it requires studying the stability to random sampling and the approximation properties of reconstruction trees. This latter result is relevant in the context of quantization of probability measures, hence of interest in its own right. We present this result first.

3.6. Approximation Error.

Towards this end, we need to introduce the infinite sample version of the reconstruction tree. For any cell $I \in \mathbb{T}$, denote the volume of the cell by

$$\rho_I = \rho(I),$$

the center of mass of the cell by

$$c_I = \begin{cases} \frac{1}{\rho_I} \int_I x \, d\rho(x) & \text{if } \rho_I > 0 \\ x_I^* & \text{if } \rho_I = 0, \end{cases}$$

where x_I^* is an arbitrary point in $\mathcal{X}$. The local expected distortion in a cell by

$$\mathcal{E}_I = \int_I \|x - c_I\|^2 \, d\rho(x),$$

and

$$\epsilon_I^2 = \mathcal{E}_I - \sum_{J\in\mathcal{C}(I)} \mathcal{E}_J = \sum_{J\in\mathcal{C}(I)} \rho_I \|c_J - c_I\|^2$$

Given the threshold $\eta > 0$, define the subtree

$$\mathcal{T}_\eta = \begin{cases} \{\mathcal{X}\} & \text{if } \epsilon_I < \eta \quad \forall I \in \mathbb{T} \\ \{I \in \mathbb{T} \mid \exists J \in \mathbb{T} \text{ such that } J \subset I \text{ and } \epsilon_J \geq \eta\} & \text{otherwise} \end{cases}, \tag{32}$$

and let $\Lambda_\eta = \Lambda(\mathcal{T}_\eta)$ be the corresponding outerleaves. Lemma 5.1.5 shows that $\mathcal{T}_\eta$ is finite, so that by Lemma 5.1.2 Λ_η is a partition and the corresponding nonlinear projection is

$$P_{\Lambda_\eta}(x) = \sum_{I \in \Lambda_\eta} c_I \mathbb{1}_I(x) \tag{33}$$

so that the code vectors are the centers of mass of the cells.

Comparing the definition of $\mathcal{T}_\eta$ and $\hat{\mathcal{T}}_\eta$, we observe that $\hat{\mathcal{T}}_\eta$ is truncated at the depth j_n given by (21), whereas $\mathcal{T}_\eta$ is not truncated, but its maximal depth is bounded by Lemma 5.1.8.

Given the above definitions, we have the following result.

THEOREM 3.6.1. *Given $\eta > 0$, for all $0 < \sigma < s$*

$$\mathcal{E}(P_{\Lambda_\eta}) \lesssim \eta^{\frac{4\sigma}{2\sigma+1}}. \tag{34}$$

Note that the bound is meaningful only if $0 < \eta < 1$. Indeed for $\eta \geq 1$ $\Lambda_\eta = \{\mathcal{X}\}$ and $\mathcal{E}(P_{\Lambda_\eta}) \leq 1$, see Remark 5.1.6.

The quantity $\mathcal{E}(P_{\Lambda_\eta})$ is called approximation error, by analogy with the corresponding definition in statistical searning theory, and is at the core of our analysis.

The problem of approximating a probability measure with a cloud of points (that is, with the corresponding counting measure) is related to the so called optimal quantization [47]. The cost of an optimal quantizer is defined as:

$$V_{N,p}(\rho) := \inf_{S \subset \mathcal{X}, |S|=N} \mathbb{E}_{x\sim\rho} d(x,S)^p$$

where $d(x,S) = \min_{y\in S} \|x-y\|$. An optimal quantizer corresponds to a set S of N points attaining the infimum, with the corresponding Voronoi-Dirichlet partition of $\mathcal{X}$. One can interpret the apporximation error $\mathcal{E}(P_{\Lambda_\eta})$ as the quantization cost associated with the (suboptimal) quantizer given by the partition Λ_η as defined in (32) with the corresponding centers $\{c_I\}_{I\in\Lambda_\eta}$, and $N := \#\Lambda_\eta$.

This point of view is also taken through the analysis of k-means given in [19], optimal quantizers corresponding in fact to absolute minimizers of the k-means problem. Asymptotic estimates for the optimal quantization cost are available, see [19] and references therein. In the special case of $\mathrm{supp}(\rho) = \mathcal{M}$, being $\mathcal{M}$ a smooth d-dimensional manifold emdedded in $\mathcal{X}$ (setting of Proposition (3.4.1)), they read:

$$\lim_{N\to\infty} N^{2/d} V_{N,2}(\rho) = C(\rho) \tag{35}$$

where $C(\rho)$ is a constant depending only on ρ. We underline that the result provided by Proposition 3.6.1 is actually a non-asymptotic estimate for the quantization cost, when the quantizer is given by the outcome of our algorithm. The quantization cost is strictly higher than the optimal one, since, for instance, an optimal quantizer always corresponds to a Voronoi-Dirichlet partition [47], whereas the partitions Λ_η are never such. Nevertheless, as observed in Section 4, a Voronoi quantizer is not suitable multiscale refinements, whereas ours is. Proposition 3.6.1 does not directly compare with (35), as it depends on a different parameter quantifying the complexity of the partition, namely η instead of N. Though, by carefully applying (55b), in the manifold case we get:

$$\mathcal{E}(P_{\Lambda_\eta}) \lesssim \left(\frac{\log N}{N}\right)^{\frac{2}{d}} \tag{36}$$

so that the bound is in fact optimal up to a logaritmic factor. Furthermore, it is in order to observe that Assumption 3.3.1 together with Proposition 3.6.1 provides a more transparent understanding of the approximation part of the analysis, as compared to other works in the literature, as we review in Section 4. It is worth stressing that Assumption 3.3.1 do not depend on the thresholding algorithm, but only on the mutual regularity of ρ and $\mathbb{T}$. Lastly we notice that, while for the sake of clarity none of the constants appear explicitly in our results, the proofs allow in principle to estimate them. Further comments about the constants are provided in Section 4.4.

4. Comparisons and Further Results

In this section a detailed comparison with the closest study in literature is provided, together with some new results and further contextualization.

4.1. Comparison with Geometric Multi-Resolution Analysis (GMRA).

A main motivation for our investigation was the algorithm GMRA [3, 63, 59], which introduces the idea of learning multi-scale dictionaries by geometric approximation.

Among the differences between GMRA and Regression Trees, the former represents data through a piece-wise linear approximation, while the latter through a piece-wise constant approximation, as outlined in Section (2.4). More precisely, rather than considering the center of mass of the data in each cell (20), a linear approximation is obtained by (local) Principal Component Analysis, so that the data belonging to a cell are orthogonally projected onto an affine subspace of suitable dimension.

Criteria of comparison might be the amount of distortion, the complexity of the representation, and the practical applicability to a specific source. Roughly speaking the piece-wise linear approach is particularly natural in the case of a ρ measure that is close to the Riemannian volume measure $\rho_{\mathcal{M}}$ of a manifold $\mathcal{M}$ of known dimension, while the piece-wise constant approach is less constraining, and might be beneficial for measures ρ that are far from evenly distributed along the support, or in case the dimension is unknown or ambiguous.

Other differences are more technical. The thresholding strategy is different: unlike [12] and our work, in [59] the local improvement $\widehat{\epsilon}$ is scaled depending on its depth in the tree. Namely, while in (22) the cells are chosen according to $\widehat{\epsilon}_I \geq \eta_n$, the GMRA thresholding is ruled by:

$$\widehat{\Delta}_{j,k} \geq 2^{-j}\tau_n$$

where $\widehat{\Delta}_{j,k}$ is formally analogous to $\widehat{\epsilon}_I$ and τ_n to η_n. See Section (2.3) in [59]. Therein this choice is justified by alleged experimental evidence as well as improved statistical guaranties. This phenomenon seems to be specific of the GMRA approach though. One of our purposes was indeed to check whether the learning rates are significantly affected in our case; as we are about to discuss, we obtain comparable rates with respect to [59].

We now provide a comparison of the analyses.

4.1.1. *Comparison of Assumptions and Results.*

In [59], Definition 5, a model class $\mathcal{B}_s$ is introduced in terms of the decay of the approximation error with respect to the scale parameter τ, following ideas from [12], where a similar model class is defined through Definition 5 therein. The problem of estimating the approximation error is essentially circumvented by grouping together all the probability measures that share the same decay. Indeed both the mentioned definitions explicitly depend on the thresholding algorithm of interest, as, for instance, Equation 7 in [59] features $\mathcal{T}_{(\rho,\eta)}$, the thresholded subtree corresponding to a certain ρ via the GRMA algorithm.

Our way to tackle the problem is substantially different; Assumption 3.3.1 describes a property of mutual regularity of the measure ρ with respect to the partition tree $\mathbb{T}$ that does not depend on the Reconstruction Tree algorithm. Based on this mutual regularity assumption, Proposition 3.6.1 addresses directly the approximation error, by providing an upper bound.

We already commented in Section 2.1 how Assumption 3.3.1 is tightly related to the so called Ahlfors regularity, and with the effective dimension of the support of ρ. Furthermore Proposition 3.4.1 provides an explicit example in which the assumption is satisfied as a consequence of simpler regularity properties.

In this vein, [59] provides Proposition 3 and Lemma 6, to address the existence of partition trees satisfying their Assumptions. In particular, the works [3, 63, 59] mentions several options as how a convenient partition tree can be built practically. The basic idea is to randomly split the dataset in two batches, one of which is used to build the partition tree $\mathbb{T}$, while the other is later utilized for the adaptive thresholding. Working with two separate batches eliminates the unwanted bias induced by cross correlation. As far as algorithms, the mentioned options are the following:

- *Cover Tree* [11]. Each Λ_j-partition of the partition tree $\mathbb{T}$ is built from a set of centroids $\mathcal{C}_j$ that is a subset of the dataset. These subsets are generated with a recursive filtering algorithm, which ensures that each $\mathcal{C}_j$ is a *cover* with respect to a radius 2^{-j}. This means that each point of the dataset is closer than 2^{-j} to some centroids of $\mathcal{C}_j$, for each j, while the centroids of $\mathcal{C}_j$ are at least 2^{-j} apart. The construction entails time complexity of order $O(C^d D n \log n)$. Once the centroids are chosen, the actual nested partitions can be built, bottom-up, by merging Voronoi cells at fine scale. See Appendix A in [59].
- *METIS* [55]. This algorithm first builds a graph by thresholding an exponential proximity kernel. The graph is then processed by a coarsening-partition-refining scheme.
- *Iterated PCA* [99]. This algorithm splits the data iteratively, by cutting $\mathcal{X}$ through an affine hyperplane, parallel to the appropriate principal components.
- *Iterated k-means* [99]. This algorithm splits the data iteratively, by performing k-means, with $k = 2$.

Among these options, the first is the one of choice in [59], by comparison of efficiency and compatibility with the assumptions of the statistical analysis. Despite the mentioned differences in our algorithm and analysis, still the Cover Tree algorithm is expected to provide an empirical counterpart of the construction we outline through the proof of Proposition 3.4.1, as the algorithm to some extent mimics the construction of dyadic cubes in [23].

Regarding our main result, Theorem 3.5.1, this bound can be compared with Theorem 8 in [59] under the assumption that $\mathcal{M}$ is a compact C^∞ manifold. Eq. (31a) with $s = 1/d$ gives a convergence rate of the order $(\ln n/n)^{\frac{2p}{2p+d}}$ for any $p = \sigma/s \in (0,1)$, whereas the GMRA algorithm has a rate of the order $(\ln n/n)^{\frac{2}{2+d}}$, see also Proposition 3 of [59]. Hence, in the limit $p \to 1$, therefore up to a logarithmic factor, our estimator has the same convergence rate of the GMRA algorithm. However it is in order to underline that our algorithm works with a cheaper representation, as long as the partitions have the same cardinality; indeed, given the adaptive partition $\hat{\Lambda}_\eta$, it only requires to compute and store the centers of mass $\{\hat{c}_I\}_{I \in \hat{\Lambda}_\eta}$.

In fact direct comparison of the rates should be taken with a grain of salt, as the cardinality of the optimal partition is specific of each algorithm. As a result, asking which representation

performs better, in terms of distortion vs compression rate, might still have a ρ-dependent answer; regardless, as already pointed out in Section 2.4, one of the representations might be more convenient for other reasons, depending on the application.

4.1.2. *The $\mathcal{A}_s$ case.*
In the mentioned papers [12] and [59] a further approximation model is introduced, in both referred to with $\mathcal{A}_s$. This model is in fact introduced before the adaptive thresholding $\mathcal{B}_s$ case, as it corresponds to the case of a *non-adaptive* truncation of the patition tree at a scale j, that is independent from the data or their distribution ρ. This method clearly provides a strictly rougher estimation, as compared to the adaptive case.

The reader might have noticed that in our analysis the *non-adaptive* model is disregarded. This is due to the fact that we decided to work with simpler assumptions (3.3.1), as compared to the mentioned papers; this assumptions, even not being specific to the algorithm at hand, still offer a way to estimate $\#\mathcal{T}_\eta$ explicitly in adaptive case (Proposition 55a); but the same machinery is not suitable to deal with the $\mathcal{A}^s$ case. Considering that the approximation obtained through a horizontally truncated subtree would be very crude in most cases, we decided to prioritize the adaptive case.

4.2. K-Means.

Our procedure being substantially a vector quantization algorithm, a comparison with the most common approach to vector quantization, namely K-means, is in order. In K-means, a set of K code vectors $c_1, \dots, c_K$ are derived from the data and used to define corresponding partitions via the corresponding Voronoi diagram

$$V_j = \{x \in \mathbb{R}^D \mid \|x - c_j\| \leq \|x - c_i\|, \forall i = 1, \dots, K, \ i \neq j\}.$$

Code vectors are defined by the minimization of the following empirical objective

$$\min_{c_1,\dots,c_K} \frac{1}{n} \sum_{i=1}^{n} \min_{j=1,\dots,K} \|x_i - c_j\|^2 .$$

This minimization problem is non convex and is typically solved by alternating minimization, a procedure referred to as Lloyd's algorithm [62]. The inner iteration assigns each point to a center, hence a corresponding Voronoi cell. The output minimization can be easily shown to update the code vectors by computing the center of mass, *the mean*, of each Voronoi cell. In general the algorithm is ensures to decrease or at least not increase the objective function and to converge in finite time to a local minimum. Clearly, the initialization is important, and initializations exist yielding some stronger convergence guarantees. In particular, K-means++ is a random initialization providing on average an ϵ approximation to the global minimum [5].
Compared to K-means, reconstruction trees restrict the search for a partition over a prescribed family defined by the partition tree. In turns, they allow a fast multi-scale exploration of the data, while K-means requires solving a new optimization problem each time K is changed (indeed it can be shown that a solution for the $K-1$ problem leads to a bad initialization for the K problem). Further, unlike restriction trees, the partitions found by K-means at different scales (different values of K) are generally unrelated, and cannot be seen one a refinement of the other.

Regarding statistical analysis, in a similar setting [19] shows that the k-means algorithm, with a suitable n-dependent choice of k, provides a convergence rate of the order $(1/n)^{\frac{1}{d+1}}$ and k-flat algorithm of order $(1/n)^{\frac{2}{d+4}}$; though these bounds might not be as tight, due to the difference in proof techniques.

4.3. Quantization via Subsampling.

In this section we argue that the learning rate given by our main result seems to be, to some extent, robust with respect to the algorithm choice. In order to show that, we present a simple algorithm, different from the one analized above, for which it's fairly easy to obtain a learning rate, at least asymptotically.

We present the infinite sample version of the algorithm first. A non-linear projection is defined, as above, by a set of N centroids $S_N = \{c_i\}_{i=1,...,N}$ and a corresponding partition $\Lambda_N = \{I_i\}_{i=1,...,N}$ of $\mathcal{X}$:

$$P_{\Lambda_N}(x) = \sum_{i=1}^{N} c_i \mathbb{1}_{I_i}(x)$$

but now we compute $S_N \sim \rho^N$ as a i.i.d. sample from ρ, and Λ_N as the Voronoi-Dirichlet partition generated by S_N.

The finite sample version of this algorithm will be defined as follows. We are given n data, that is, a i.i.d. sample $\{x_1, ..., x_n\} \sim \rho^n$. The centroids $\widehat{S}_N$, with $N \leq n$, will be chosen by uniform subsampling without replacement from $\{x_1, ..., x_n\}$; namely, $\widehat{S}_N$ is obtain from $\widehat{S}_{N-1}$ by drawing a new centroid from $\{x_1, ..., x_n\} \setminus \widehat{S}_{N-1}$ with uniform probability. Notice that every subset of N variables $\{x_{i_1}, x_{i_N}\}$ is drawn from the same measure ρ^N, and each of them has the same probability to be chosen as $\widehat{S}_N$. So $\widehat{S}_N \sim \rho^N$.

More precisely, let I be the uniform distribution across $\{1, \ldots, n\}$, then the sampling of indexes $I_1, \ldots, I_N$ without replacement implies that these indexes are distributed according to the corresponding multihypergeometric distribution. Furthermore $I_1, \ldots, I_N$ are independent from $\{x_1, ..., x_n\}$. Now let $y_1, \ldots, y_n$ be such that $y_i = x_{I_i}$. The claim corresponds to saying that $y_i, \ldots, y_n$ are identically and independently distributed according to ρ.

Let $E \subseteq \mathbb{R}^D$. Then:

$$\mathbb{P}[x_{I_i} \in E] = \sum_{j=1}^{n} \mathbb{P}[x_j \in E | I_i = j] \mathbb{P}[I_i = j],$$

where $I_i \perp x_j$, so that:

$$= \sum_{j=1}^{n} \mathbb{P}[x_j \in E] \mathbb{P}[I_i = j] = \mathbb{P}[x_j \in E] \sum_{j=1}^{n} \mathbb{P}[I_i = j] = \mathbb{P}[x_j \in E],$$

therefore $y_i = x_{I_i} \sim \rho$. Now:

$$\mathbb{P}[x_{I_{i_1}} \in E_1, \ldots, x_{I_{i_k}} \in E_k] = \sum_{j_1,...,j_k=1}^{n} \mathbb{P}[x_{j_1} \in E_1, \ldots, x_{j_k} \in E_k | I_i = j] \mathbb{P}[I_{j_1} = j_1, \ldots, I_{j_k} = j_k],$$

and since $j_1 \neq \cdots \neq j_k$, so that:

$$= \mathbb{P}[x_{j_1} \in E_1] \cdot \cdots \cdot \mathbb{P}[x_{j_k} \in E_k | I_i = j] \sum_{j_1,...,j_k=1}^{n} \mathbb{P}[I_{j_1} = j_1, \ldots, I_{j_k} = j_k]$$

and the claim follows.

Therefore there is no actual difference between $\widehat{S}_N$ and S_N, which is the same as saying that $\widehat{S}_N$ has no extra variance due to the size of the sample n; S_N itself does have variance though, depending on N, because the algorithm is randomized. Notice that $P_{\Lambda_N}\rho = \sum_I \rho(I)\delta_{c_I}$; this

differs from what is often called empirical measure, namely $\widehat{\rho}_N = \sum_{I\in\Lambda_N} \frac{1}{N}\delta_{c_I}$. By applying Lemma 3.1 from [20] we obtain that:

$$\mathcal{E}(P_{\Lambda_N}) = W_2^2(\rho, P_{\Lambda_N}\rho) = \Theta\left(N^{-\frac{2}{d}}\right) \tag{37}$$

where $W_2(,)$ is the 2-Wasserstein distance between probability measures, and the second equality is discussed in [19, 20].

To compare (37) with Corollary 3.5.2 in the manifold case, we proceed as follows. We impose the value of N_n to be comparable with $\#\Lambda_{\eta_n}$, so that the algorithms produce comparable outcomes. This is done by using the estimates (30) and (55b), namely $N_n \propto (\frac{n}{\log n})^{\frac{d}{d+2}}$, so that the (37) becomes:

$$\mathcal{E}(P_{N_n}) = \Theta\left(\left(\frac{\log n}{n}\right)^{\frac{2}{d+2}}\right) \tag{38}$$

so that, asymptotically, the uniform subsampling algorithm achieves the same learning rate as in our main result.

Lastly we make the reader notice that, as far as multiscale refinements are concerned, the uniform subsampling algorithm allows in fact to be performed coarse-to-fine. Indeed, as described above, S_N is built one centroid at a time, and Λ_N is computed, as Voronoi-Dirichlet partition of S_N, from Λ_{N-1} by updating the only cells corresponding to centroids of S_{N-1} that are nearest neighboors of the new centroid in S_N.

4.4. Pseudo-Normas.

This section contains a remark that arose during our investigation, by comparing our result and our main sources; this is at the moment still at an early stage of understanding.

We invite the reader to compare the following three constants:

(1) Consider the quantity $|f|_{\mathcal{B}^s}$ in [12] at (2.1.6)

$$\|f - P_{\Lambda_{(f,\eta)}}(f)\| \leq C_s |f|_{\mathcal{B}^s} (\#\mathcal{T}(f,\eta))^{-s}$$

where f is a function being learned via adaptive non-linear projection $P_{\Lambda_{(f,\eta)}}(f)$, and C_s is a constant depending only on s.

(2) We already mentioned (35), from [19]. Its extended form reads:

$$\lim_{N\to\infty} N^{2/d} V_{N,2}(\rho) = C_{2,d}\left(\int_{\mathcal{M}} d\mu_{\mathrm{I}} p(x)^{\frac{d}{d+2}}\right)^{\frac{d+2}{d}}$$

where $C_{2,d}$ is a constant depending only on d, $d\mu_{\mathrm{I}}$ is the first fundamental form of $\mathcal{M}$, and p is the density associated with the probability measure ρ; in this case it is assumed that such a density is well defined, as in Proposition 3.4.1.

(3) By inspecting the proof of (55a), in the form with s rather than σ, we observe that the constant, implicit in $\lesssim$, only depends on the constants C, b from the Assumptions. In fact the following lemma provides an explicit form for the constant, under assumptions along the line of Proposition 3.4.1.

LEMMA 4.4.1. *Assume that ρ has density p with respect to a probability measure μ and the following two assumptions hold true*

a) there are two constants $\Phi_-, \Phi_+ > 0$ such that for all $I \in \Lambda$ with $\rho_I > 0$

$$\Phi_- \, \rho_I \leq p(x) \, \mu_I \leq \Phi_+ \, \rho_I \qquad \mu\text{-almost all } x \in I \qquad \Phi_- \, \rho_I \leq \mu_I \tag{39}$$

where $\mu_I = \mu(I)$;

b) there exists $s \in (0, +\infty)$ and $D_s > 0$ such that

$$\operatorname{diam}_\mu I \leq D_s \, \mu_I^s \qquad I \in \Lambda \tag{40}$$

then

$$\begin{aligned} \sharp\{I \in \Lambda : E_I \geq \eta^2\} &\leq \Phi_-^{-1} \left(D_s^2 \, \Phi_+ \, \|p\|_{\frac{1}{2s+1}} \right)^{\frac{1}{2s+1}} \eta^{-\frac{1}{2s+1}} \\ &\leq \Phi_-^{-1} \left(D_s^2 \, \Phi_+ \left\| p^{-2s} \mathbb{1}_{|p(x)>0} \right\| \right)^{\frac{1}{2s+1}} \eta^{-\frac{1}{2s+1}}. \end{aligned} \tag{41}$$

The first inequality in (39) implies $p(x) > 0$ μ-almost everywhere, so that the measures ρ and μ restricted to I are equivalent. The blue inequality in (39) is an alternative independent condition with respect to the first inequality in (39). A typical example is when $\mathcal{M}$ is a manifold with dimension d and μ is the Riemannian volume element of $\mathcal{M}$. If $I \cap \mathcal{M}$ are *ball-like* , then we expect that $s = 1/d$.

Recalling that the constant $|f|_{\mathcal{B}^s}$ is directly related to the Besov norm $\|\cdot\|_{\mathcal{B}^s_q}$ in the case of classical wavelets for function approximation (see [24]), it turns out that a direct comparison of the constants in the three instances of least square approximation error above, that is, form [12], [19] and the bound (36) together with Lemma 4.4.1, suggest that the constant expressed via $\left(\int_{\mathcal{M}} d\mu_I p(x)^{\frac{d}{d+2}} \right)^{\frac{d+2}{d}}$ might provide an analogue of the Besov norm in the manifold setting of the present work. This is at the moment just a rough idea.

Proof of Lemma 4.4.1 is provided in the next section.

4.5. Further Comparisons.

4.5.1. *Decision and Regression Trees.* The above procedure was named Reconstruction Tree because of its formal analogy to that of Decision Trees for supervised learning; see for example [54], Chapter 8. In fact, as already mentioned, our construction and analysis follows closely that of tree based estimators studied in [12], in the context of least square regression, whose estimator can be Regression Tree accordingly. Despite the formal similarity, the two settings do exhibit distinct features. For example, the analysis in [12] is specifically formulated for scalar functions, while our analysis is necessarily vectorial in nature. In [12] a uniform bound $|y| < M$ is imposed, while in our setting we can assume a local bound for free; namely, if f is constant on a cell $I \subset \mathcal{X}$ then $\|x - f(x)\|^2 \leq \operatorname{diam}(I)^2$ for all $x \in I$. The present setting finds a natural instance in the case of a probability measure supported on a smooth manifold isometrically embedded in $\mathcal{X}$ (see Section 3.3), while this case is hardly addressed explicitly in the literature about least square regression. On the other hand, the manifold case is actually discussed, in the context of classification through Decision Trees, in [93].

4.5.2. *Empirical Risk Minimization.* Again in analogy to supervised learning, as in [12], one can consider the minimization problem:

$$\min_{F \in \mathcal{H}} \widehat{\mathcal{E}}[F],$$

where $\mathcal{H}$ is the (finite-dimensional) vector space of the vector fields $F : \mathcal{X} \to \mathbb{R}^D$, which are piece-wise constant on a given partition Λ. There correspond a number of independent minimization problems, one for each cell in Λ, so that $\widehat{P}_\eta$ from (23) is easily shown to be a minimizer. The minimizer is not unique, since the value of F is irrelevant on cells $I \in \Lambda$ such that $n_I = 0$. Similar considerations hold for $\min_{F \in \mathcal{H}} \mathcal{E}[F]$ as well, in which case the value of F is irrelevant whenever $\rho_I = 0$. See also Section 3.2, Section 3.6 and Lemma 7.1.2.

One could consider minimization over a wider class of functions, piece-wise constant on different partitions, for example on all the partitions with a given number of cells that are induced by proper subtrees $\mathcal{T}$ of a given partition tree $\mathbb{T}$. This would be a combinatorial optimization problem. The algorithm defined in (22) overcomes this issue by providing a one-parameter coarse-to-fine class of partitions, such that each refinement carries local improvements $\widehat{\epsilon}$ that are uniformly bounded. As observed in [12], such a strategy is inspired by wavelet thresholding.

4.5.3. *Hierarchical Clustering.* Lastly, our coarse-to-fine approach can be compared with hierarchical clustering, in particular with the so called Ward's method, which proceeds in the opposite way. Indeed, this algorithm produces a coarser partition of the data starting from a finer. It starts with a Voronoi partition having all the data as centers, and at each step it merges a couple of cells that have the smallest so called between cluster inertia [108]. Interestingly this definition has an analogue in our algorithm. Our $\widehat{\mathcal{E}}_I$ corresponds to the within cluster inertia of a cell I while $\widehat{\epsilon}_I^2$ to the between cluster inertia (up to a factor $1/n$) of cells that merge into I. Nevertheless the obtained partitions will not in general coincide.

5. Proofs of Approximation Results

Structural results about the partition tree and how its size is estimated are presented first. Statistical results and stability properties are delayed to later sections.

5.1. Structure of the partition tree and approximation. We introduce some further notations about partition trees.

With slight abuse of notation, we regard $\mathbb{T} = \cup_{j\in\mathbb{N}} \cup_{I\in\Lambda_j} I$ as the (disjoint) union of the cells in each partition Λ_j. It a cell does not split, *i.e.* $\mathcal{C}(I) = I$, we regard $I \in \Lambda_j$ and $\mathcal{C}(I) \in \Lambda_{j+1}$ as different cells.

Given a cell $I \in \mathbb{T}$, for any $k \in \mathbb{N}$ we set

$$\mathcal{C}^{k+1}(I) = \mathcal{C}(\mathcal{C}^k(I)) \qquad \mathcal{P}^{k+1}(I) = \mathcal{P}(\mathcal{P}^k(I)),$$

where $\mathcal{C}^0(I) = \mathcal{P}^0(I) = \{I\}$ and, clearly,

$$\mathcal{P}^k(I) = \{\mathcal{X}\} \qquad k \geq j_I.$$

Furthermore, for any pair $I, J \in \mathbb{T}$, $I \neq J$ one and only one of the following alternative possibilities holds true

$$I \cap J = \varnothing \qquad \text{or} \qquad J \in \mathcal{C}^k(I) \qquad \text{or} \qquad I \in \mathcal{C}^k(J) \tag{42}$$

for some $k \geq 1$.

If $\{\mathcal{T}_t\}_{t\in T}$ is an arbitrary family of subtrees, clearly the intersection $\cap_{t\in T}\mathcal{T}_t$ and the union $\cup_t \mathcal{T}_t$ are the *smallest* and the *largest* subtrees in the family.

Given a subset $S \subset \mathbb{T}$, we set

$$\mathcal{T}(S) = \bigcap \{\mathcal{T} \mid \mathcal{T} \text{ is a subtree and } S \subset \mathcal{T}\},$$

which is the smallest subtree containing all the cells in S, and

$$\mathcal{T}(S) = \bigcup_{I\in S} \{\mathcal{P}^k(I) \mid k = 0, \dots, j_I\} = \bigcup_{I\in S} \mathcal{T}(I) \tag{43a}$$

$$\sharp\mathcal{T}(S) \leq 1 + \sum_{I\in S} j_I. \tag{43b}$$

The following remark provides an alternative definition of $\hat{\mathcal{T}}_\eta$ and a similar procedure can be applied to $\mathcal{T}_\eta$

Remark 5.1.1. *Given $\eta > 0$ and $j_n \in \mathbb{N}$, set*

$$\hat{S}_\eta = \{I \in \mathbb{T} \mid j_I \geq j_n, \hat{\epsilon}_I \geq \eta\}$$

which is a finite subset of $\mathbb{T}$. It is easy to check that $\mathcal{T}(\hat{S}_\eta) = \hat{T}_\eta$, see Figure 1.

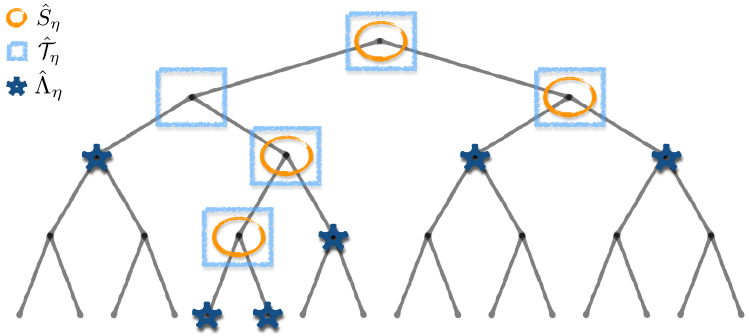

FIGURE 1. The adaptive tree

The following lemma is quite obvious.

LEMMA 5.1.2. *If $\mathcal{T}$ is finite subtree, the family of outer leaves*

$$\Lambda(\mathcal{T}) = \{I \in \mathbb{T} : I \notin \mathcal{T}, \mathcal{P}(I) \in \mathcal{T}\}$$

is a partition with

$$\sharp\Lambda(\mathcal{T}) \leq (a-1)\,\sharp\mathcal{T} + 1 \leq a\,\sharp\mathcal{T}. \tag{44}$$

PROOF. For any $j \in \mathbb{N}$, we fix an arbitrary order among the cells in Λ_j, *i.e.*

$$\Lambda_j = \{I_{j,1}, \dots, I_{j,N_j}\} \qquad N_j = \#\Lambda_j.$$

Given two different cells $I = I_{j,k}$ and $J = I_{j',k'}$ in $\mathbb{T}$, we say that I is older than J either if $j < j'$ or $j = j'$ and $k < k'$.

By definition $\mathcal{T}$ contains the parents of all its elements and, hence, the root $\mathcal{X}$. Define, by induction, the family of subtrees

$$\mathcal{T}_1 = \{\mathcal{X}\} \qquad \mathcal{T}_{n+1} = \mathcal{T}_n \cup \{I\},$$

where I is the oldest cell in $\mathcal{T} \setminus \mathcal{T}_n$. Note that, by construction, $I \in \Lambda(\mathcal{T}_n)$. Since $\mathcal{T}$ is finite by assumption, $\mathcal{T}_N = \mathcal{T}$ with $N = \sharp\mathcal{T}$.

We now prove by induction on $n = 1, \dots, N$ that $\Lambda(\mathcal{T}_n)$ is a partition. If $n = 1$

$$\Lambda(\mathcal{T}_1) = \{I \in \mathbb{T} | I \neq \mathcal{X}, \mathcal{P}(I) = \mathcal{X}\} = \mathcal{C}(\mathcal{X}) = \Lambda_1$$

which is a partition by assumption and, by (19), it satisfies (44). Assume that $\Lambda(\mathcal{T}_n)$ is partition satisfying (44) with $\mathcal{T} = \mathcal{T}_n$. By construction, $\mathcal{T}_{n+1} = \mathcal{T}_n \cup I$ with $I \in \Lambda(\mathcal{T}_n)$, then

$$\Lambda(\mathcal{T}_{n+1}) = \left(\bigcup_{J \in C(I)} \{J\} \right) \bigcup (\Lambda(\mathcal{T}_n) \setminus \{I\})$$

which is a partition since $I = \cup_{J \in C(I)} J$ and

$$\sharp\Lambda(\mathcal{T}_{n+1}) \leq a + \sharp\Lambda(\mathcal{T}_n) - 1 \leq a + (a-1)\sharp\mathcal{T}_n = (a-1)(\sharp\mathcal{T}_n + 1) + 1 = (a-1)(\sharp\mathcal{T}_{n+1}) + 1,$$

so that (44) holds true with $\mathcal{T} = \mathcal{T}_{n+1}$. □

We observe that, given a cell I, by definition of $\mathrm{diam}_\rho(I)$, it exists a measurable subset $J \subset I$ such that $\rho(J) = 0$ and $\mathrm{diam}(I \setminus J) = \mathrm{diam}_\rho(I)$. Furthermore,

$$\mathbb{P}\left[\|X_i - X_j\| > \mathrm{diam}_\rho(I) | X_i, X_j \in I\right] = 0 \qquad i, j = 1, \dots, n.$$

The following simple lemmas will be useful.

LEMMA 5.1.3. *Given a cell $I \in \mathbb{T}$ with $\rho(I) > 0$*

$$\|x - c_I\| \leq \mathrm{diam}_\rho(I) \qquad \rho - \textit{almost all } x \in I \tag{45a}$$

$$\|c_J - c_K\| \leq \mathrm{diam}_\rho(I) \qquad J, K \in \mathcal{C}(I) \tag{45b}$$

$$\|\widehat{c}_I - c_I\| \leq \begin{cases} \mathrm{diam}_\rho(I) & n_I \neq 0 \\ \mathrm{diam}(\mathcal{X}) & n_I = 0 \end{cases} \qquad \textit{almost surely} \tag{45c}$$

PROOF. By definition of essential diameter, there exists $I_0 \subset I$ such that $\mathrm{diam}(I_0) = \mathrm{diam}_\rho(I)$ and $\rho(I \setminus I_0) = 0$. Let C the closed convex hull of I_0. It is known that $\mathrm{diam}(C) = \mathrm{diam}_\rho(I)$ and, by convexity theorem, see [92, Thm. 5.7.35],

$$c_I = \frac{1}{\rho(I_0)} \int_{I_0} x \, d\rho(x) \in C,$$

so that (45a) is clear. Since $J, K \subset I$, Eq. (45b) is a consequence of the fact that $c_J, c_K \in C$. If $n_I = 0$, $\widehat{c}_I = \hat{x}_I^* \in \mathcal{X}$ so that (45c) is clear. If $n_I \neq 0$, almost surely $\widehat{c}_I \in C$ so that

$$\|\widehat{c}_I - c_I\| \leq \mathrm{diam}(C) = \mathrm{diam}(I_0) = \mathrm{diam}_\rho(I).$$

□

Given a cell $I \in \mathbb{T}$, the within-between decomposition of the variance

$$\mathcal{E}_I = \sum_{J \in \mathcal{C}(I)} \mathcal{E}_J + \sum_{J \in \mathcal{C}(I)} \rho_I \|c_J - c_I\|^2, \tag{46}$$

implies

$$\epsilon_I^2 = \sum_{J \in \mathcal{C}(I)} \rho_I \|c_J - c_I\|^2 \tag{47}$$

As a consequence we have the following decomposition.

LEMMA 5.1.4. *Given $I \in \mathbb{T}$, for all $N \in \mathbb{N}$*

$$\mathcal{E}_I = \sum_{k=0}^{N} \sum_{J \in \mathcal{C}^k(I)} \epsilon_J^2 + \sum_{J \in \mathcal{C}^{N+1}(I)} \mathcal{E}_J. \tag{48}$$

PROOF. The claim is clear for $N = 0$. Assume that it holds true for N. Then, for any $J \in \mathcal{C}^{N+1}(I)$

$$\mathcal{E}_J = \epsilon_J^2 + \sum_{J' \in \mathcal{C}(J)} \mathcal{E}_{J'},$$

hence

$$\mathcal{E}_I = \sum_{k=0}^{N} \sum_{J \in \mathcal{C}^k(I)} \epsilon_J^2 + \sum_{J \in \mathcal{C}^{N+1}(I)} \left(\epsilon_J^2 + \sum_{J' \in \mathcal{C}(J)} \mathcal{E}_{J'} \right),$$

by observing that a cell $J' \in \mathcal{C}(J)$ for some $J \in \mathcal{C}^{N+1}(I)$ if and only if $J' \in \mathcal{C}^{N+2}(I)$, so that (48) holds true for $N+1$. □

We now show that the set $\mathcal{T}_\eta$ defined by (32) is a finite subtree.

LEMMA 5.1.5. *The family $\mathcal{T}_\eta$ is a finite subtree of $\mathbb{T}$.*

PROOF. If $\mathcal{T}_\eta = \{\mathcal{X}\}$, there is nothing to prove. Otherwise, if $I \in \mathcal{T}_\eta$, then by definition there exists $J \in \mathbb{T}$ such that $J \subset I$ and $\epsilon_J \geq \eta$. Since $\mathcal{P}(I) \supset I \supset J$, then $P(I) \in \mathcal{T}_\eta$, so that $\mathcal{T}_\eta$ is a subtree.

We now show that $\mathcal{T}_\eta$ is finite. From (48) with $\mathcal{X} = I$ and (50) we get that for all $N \in \mathbb{N}$

$$\sum_{k=0}^{N} \sum_{J \in \mathcal{C}^k(\mathcal{X})} \epsilon_J^2 \leq \mathcal{E}_\mathcal{X} \leq 1.$$

Then the series $\sum_{k=0}^{+\infty} \sum_{J \in \mathcal{C}^k(\mathcal{X})} \epsilon_J^2 = \sum_{I \in \mathbb{T}} \epsilon_J^2$ is sommable. Hence, the set

$$S_\eta = \{I \in \mathbb{T} : \epsilon_I \geq \eta, \}$$

is finite. Furthermore, by construction

$$\mathcal{T}_\eta = \mathcal{T}(S_\eta) = \bigcup_{I \in S_\eta} \mathcal{T}(I). \tag{49}$$

see Figure 1. Bound (43b) implies that $\mathcal{T}_\eta$ is finite. □

REMARK 5.1.6. *By* (24),

$$\epsilon_I^2 \leq \mathcal{E}_I \leq \operatorname{diam}_\rho(I)^2 \rho_I \leq \operatorname{diam}(\mathcal{X})^2 \leq 1, \tag{50}$$

and $\epsilon_I < \mathcal{E}_I \leq 1$ provided that $\epsilon_I > 0$. Furthermore, by (50) *$\mathcal{T}_\eta = \Lambda_\eta = \{\mathcal{X}\}$ for all $\eta \geq 1$ and $\mathcal{E}(P_{\Lambda_\eta}) = \mathcal{E}_\mathcal{X} \leq 1$. By the same argument $\hat{\mathcal{T}}_\eta = \hat{\Lambda}_\eta = \{\mathcal{X}\}$. Hence, it is enough to consider the case $0 < \eta < 1$.*

5.1.1. *A-term: Approximation error.* In this section, we bound the approximation error, which is based in an estimation of the number of cells whose loss of inertia is big enough.

LEMMA 5.1.7. *Given a partition $\Lambda \subset \mathbb{T}$, given $0 < \eta < 1$,*

$$\sharp\{I \in \Lambda : \epsilon_I \geq \eta\} \leq \sharp\{I \in \Lambda : \mathcal{E}_I \geq \eta^2\} \lesssim \eta^{-\frac{2}{2s+1}}. \tag{51}$$

PROOF. First inequality in (51) is a direct consequence of (50). By (25a) there exists $C > 0$ such that for all $I \in \mathbb{T}$

$$\operatorname{diam}_\rho(I) \leq C \rho_I^s \qquad I \in \Lambda,$$

then, by (50)

$$\mathcal{E}_I \leq C^2 \rho_I^{2s+1}.$$

Set $\Lambda_+ = \{I \in \Lambda : \mathcal{E}_I \geq \eta^2\}$ and $N_+ = \sharp\Lambda_+$. Fix $q \geq 1$, clearly

$$\eta^{\frac{2}{q}} N_+ \leq \sum_{I \in \Lambda_+} \mathcal{E}_I^{\frac{1}{q}} \leq C^{\frac{2}{q}} \sum_{I \in \Lambda_+} \rho_I^{\frac{2s+1}{q}}. \tag{52}$$

The Holder inequality with $1/p + 1/q = 1$ gives

$$\sum_{I\in\Lambda_+} \rho_I^{\frac{2s+1}{q}} \leq \left(\sum_{I\in\Lambda_+} \rho_I^{\frac{p(2s+1)}{q}}\right)^{\frac{1}{p}} \left(\sum_{I\in\Lambda_+} 1^q\right)^{\frac{1}{q}} \leq \left(\sum_{I\in\Lambda} \rho_I\right)^{\frac{2s+1}{2s+2}} N_+^{\frac{1}{2s+2}} = N_+^{\frac{1}{2s+2}}$$

where the last inequality follows by choosing $\frac{p(2s+1)}{q} = 1$, *i.e.* $q = 2s+2$. By replacing in (52)

$$N_+^{1-\frac{1}{2s+2}} \leq C^{\frac{1}{s+1}} \eta^{-\frac{2}{2s+2}}$$

and, since $1 - \frac{1}{2s+2} = \frac{2s+1}{2s+2}$, we get (51). □

We first observe that the proof of the above lemma only depends on Assumption (25a) and the constant in the inequality (51) only depends on the constant in (25a), denoted by C in the proof. Furthermore, without Assumption (25a) we always have the following bound

$$\sharp\{I \in \Lambda : \mathcal{E}_I \geq \eta^2\} \leq \sharp\{I \in \Lambda : \mathcal{E}_I \geq \eta^2\} \leq \eta^{-2}.$$

where the first inequality is consequence of (50) and the last bounds is due to (46) with $I = \mathcal{X}$

$$1 \geq \mathcal{E}_{\mathcal{X}} \geq \sum_{I\in\Lambda} \mathcal{E}_I \geq \sum_{I\in\Lambda, \mathcal{E}_I\geq\eta^2} \mathcal{E}_I = \eta^2 \sharp\{I \in \Lambda : \mathcal{E}_I \geq \eta^2\}.$$

We recall that $\Lambda_\eta = \Lambda(T_\eta)$ is the family of the corresponding outer leaves, which is a partition of $\mathcal{X}$ by Lemma 5.1.2. In order to bound the cardinality of Λ_η we need an auxiliary lemma based on Assumption (25b).

LEMMA 5.1.8. *Given $\eta > 0$, set*

$$j_\eta = \sup\{j_I \in \mathbb{N} \mid I \in \mathbb{T} \text{ and } \epsilon_I \geq \eta\}. \tag{53}$$

then

$$j_\eta \lesssim \ln(\frac{2}{\eta}). \tag{54}$$

If $\{j_I \in \mathbb{N} \mid I \in \mathbb{T} \text{ and } \epsilon_I \geq \eta\} = \varnothing$ we set $j_\eta = 0$.

PROOF. If $\mathcal{T}_\eta = \{\mathcal{X}\}$ or $\{j_I \in \mathbb{N} \mid I \in \mathbb{T} \text{ and } \epsilon_I \geq \eta\} = \varnothing$, then $j_\eta = 0$, so that the claim is evident. If $\mathcal{T}_\eta \neq \{\mathcal{X}\}$, then $0 < \eta < 1$. Take $I \in \mathbb{T}$ such that $\epsilon_I \geq \eta$. By (50) and (25b)

$$\eta^2 \leq \epsilon_I^2 \leq \rho_I \operatorname{diam}_\rho(I)^2 \leq C_2 b^{-2j_I}$$

Hence

$$j_I \leq \frac{1}{\ln b} \ln(\frac{1}{\eta}) + \frac{1}{2\ln b} \ln C_2 \leq E \ln(\frac{2}{\eta}),$$

where $E = \max\{1, \frac{\ln C_2}{2\ln 2}\}/\ln b$. □

PROPOSITION 5.1.9. *Given $\eta > 0$, for all $\sigma < s$,*

$$\sharp\mathcal{T}_\eta \lesssim \eta^{-\frac{2}{2s+1}} \ln(\frac{2}{\eta}) \lesssim \eta^{-\frac{2}{2\sigma+1}} \tag{55a}$$

$$\sharp\Lambda_\eta \lesssim \eta^{-\frac{2}{2s+1}} \ln(\frac{2}{\eta}) \lesssim \eta^{-\frac{2}{2\sigma+1}}, \tag{55b}$$

where the constants in $\lesssim$ also depend on σ.

PROOF. As observed in Remark 5.1.6, we can assume that $0 < \eta < 1$. Let

$$\Upsilon = \{I \in \mathcal{T}_\eta \mid \epsilon_I \geq \eta \text{ and } \epsilon_J < \eta \, \forall J \in \mathcal{C}^k(I),\ k \geq 1\}.$$

By (42) the elements of Υ are disjoint. Let $\Lambda \subset \mathbb{T}$ be a partition such that $\Upsilon \subset \Lambda$. Hence, by (51)

$$\sharp\Upsilon \leq \sharp\{I \in \Lambda \mid \epsilon_I \geq \eta\} \lesssim \eta^{-\frac{2}{2s+1}}. \tag{56}$$

We claim that

$$\mathcal{T}_\eta = \bigcup_{J \in \Upsilon} \mathcal{T}(J) = \mathcal{T}(\Upsilon).$$

By construction $\mathcal{T}_\eta \supset \bigcup_{J\in\Upsilon} \mathcal{T}(J)$. To prove the opposite inclusion, fix $I \in \mathcal{T}_\eta$, then there exists $J_1 \in \mathbb{T}$ such that $J_1 \subset I$ and $\epsilon_{J_1} \geq \eta$. If $J_1 \in \Upsilon$, then $I \in \mathcal{T}(J_1)$. Otherwise, there exists $J_2 \subset J_1 \subset I$ and $\epsilon_{J_2} \geq \eta$. If $J_2 \in \Upsilon$, $I \in \mathcal{T}(J_2)$. Otherwise, because of $\mathcal{T}_\eta$ is finite, we can repeat the procedure until we get $J_k \in \Upsilon$ such that $J_k \subset I$, then $I \in \mathcal{T}(J_k)$ and the claim is proven. By (43b)

$$\sharp\mathcal{T}_\eta = \sharp\mathcal{T}(\Upsilon) \leq 1 + \sum_{J\in\Upsilon} j_J \leq 1 + \sharp\Upsilon\ j_\eta \lesssim \eta^{-\frac{2}{2s+1}} \ln(\frac{2}{\eta}),$$

since, by definition, $j_J \leq j_\eta$ and the last inequality is due to (56) and (54). This shows the first inequality in (55a). Since $\sigma < s$, for some $\delta > 0$

$$\eta^{-\frac{2}{2s+1}} \ln(\frac{2}{\eta}) = \eta^{-\frac{2}{2\sigma+1}} \eta^\delta \ln(\frac{2}{\eta}) \leq C\eta^{-\frac{2}{2\sigma+1}}$$

where $C = \sup_{0<\eta\leq 1} \eta^\delta \ln(\frac{2}{\eta})$, which is finite, since $\lim_{\eta\to 0} \eta^\delta \ln(\frac{2}{\eta}) = 0$. Bound (55b) is a direct consequence of (44). □

REMARK 5.1.10. *In the following, for sake of clarity we bound the logarithmic dependence on η by considering $\sigma < s$. However our results can be extended to $\sigma = s$ by adding a logarithmic factor, as in* (55a) *and* (55b).

5.2. Manifold Setting Results.

PROOF OF THM. 3.4.1. We first observe that it is enough to show that there exists a partition tree $\mathbb{T}' = \{\Lambda_j\}$ for $\mathcal{M}$. Indeed, by adding to each partion Λ_j, the cell $I_0 = \mathcal{X} \setminus \mathcal{M}$, we get a partion of $\mathcal{X}$, which satisfies Assumptions 3.3.1, since $\operatorname{diam}_\rho(I_0) = 0$.

Since $\mathcal{X}$ is bounded, then $\mathcal{M}$ is a connected compact manifold and, hence, $(\mathcal{M}, d_\mathcal{M}, \rho_\mathcal{M})$ is an Ahlfors regular metric measure space [46, page 413], *i.e.*

$$d_1\, r^d \leq \rho_\mathcal{M}(B_\mathcal{M}(x,r)) \leq d_2\, r^d \qquad r \leq \operatorname{diam}(\mathcal{M}),$$

where $B_\mathcal{M}(x,r)$ is the ball of center x and radius r with respect to the Riemannian metric $d_\mathcal{M}$. By (28b)

$$d_1 p_1\, r^d \leq \rho(B_\mathcal{M}(x,r)) \leq d_2 p_2\, r^d \qquad r \leq \operatorname{diam}(\mathcal{M}), \tag{57}$$

where d is the intrinsic dimension of $\mathcal{M}$. Since $(\mathcal{M}, d_\mathcal{M}, \rho)$ is an Ahlfors regular metric measure, too, there exists a family of dyadic cubes, *i.e* for each $j \in \mathbb{Z}$ there is a family $\Lambda_j = \{I\}$ of open subsets of $\mathcal{M}$ such that

$$\rho(\mathcal{M} \setminus \cup_{I\in\Lambda_j} I) = 0 \tag{58a}$$

$$I \cap J = \varnothing \qquad I, J \in \Lambda_j,\ I \neq J \tag{58b}$$

$$\text{either } I \cap J = \varnothing \text{ or } J \subset I \qquad I \in \Lambda_j,\ J \in \Lambda_{j+\ell} \tag{58c}$$

$$I \supset B_{\mathcal{M}}(x_I, r_0\delta^j) \qquad I \in \Lambda_j \tag{58d}$$

$$I \subset B_{\mathcal{M}}(x_I, r_1\delta^j) \qquad I \in \Lambda_j \tag{58e}$$

where $0 < r_0 < r_1$ and $\delta \in (0,1)$ are given constants [23, Theorem 11]. As noted in [42], it is always possible to redefine each cell $I \in \Lambda_j$ by adding a suitable portion of its boundary in such a way that

$$\mathcal{M} = \cup_{I\in\Lambda_j} I \tag{59}$$

and (28a)–(58e) still hold true, possibly with different constants. Since $\mathcal{M}$ is compact, there exists $j_0 \in \mathbb{Z}$ such that $B_{\mathcal{M}}(x_0, r_1\delta^{j_0}) = \mathcal{M}$ for some $x_0 \in \mathcal{M}$. Hence, possibly redefining j, r_0 and r_1, we can assume that $\Lambda_0 = \{\mathcal{M}\}$ and, as a consequence of (58b), (58c) and (59), the family $\{\Lambda_j\}_{j\in\mathbb{N}}$ is a partion tree for $\mathcal{M}$ where the bound in (19) is a consequence of the following standard volume argument. Fix j_0 large enough such that for all $j \geq j_0$, $r_1\delta^j \leq \operatorname{diam}\mathcal{M}$, then given $I \in \Lambda_j$

$$\rho(I) = \sum_{J\in\mathcal{C}(I)} \rho(J) \geq \sum_{J\in\mathcal{C}(I)} \rho(B_{\mathcal{M}}(x_J, r_0\delta^{j+1})) \geq \sharp\mathcal{C}(I)\, d_1 p_1 r_0^d \delta^{d(j+1)},$$

where the third and the forth inequalities are consequence of (58d) and (57).

On the other hand, by (58e) and (57),

$$\rho(I) \leq \rho(B_{\mathcal{M}}(x_I, r_1\delta^j)) \leq d_2 p_2 r_1^d \delta^{jd}$$

so that

$$\sharp\mathcal{C}(I) \leq \frac{d_2 p_2 r_1^d}{d_1 p_1 r_0^d \delta^d} = D$$

Bound (19) holds true by setting

$$a = \max\{\max_{\substack{j<j_0\\ I\in\Lambda_J}} \{\#\,\mathcal{C}(I)\}, D\}.$$

We now show that (25b) holds true. Indeed, since $\mathcal{M}$ is Riemmannian submanifold of $\mathbb{R}^D$ it holds that

$$\|y - x\| \leq d_{\mathcal{M}}(y, x) \qquad x, y \in \mathcal{M}, \tag{60}$$

see [82, Corollary 2 and Proposition 21, Chapter 5]. Given $I \in \Lambda_j$, by (58e),

$$\operatorname{diam}_\rho(I) \leq \operatorname{diam}(I) \leq \sup_{x,y\in I} \|x - y\| \leq \sup_{x,y\in B_{\mathcal{M}}(x_I, r_1\delta^j)} d_{\mathcal{M}}(x, y) \leq 2r_1\delta^j,$$

so that (25b) holds true with $b = 1/\delta > 1$ and $C_2 = 2r_1$. To show (25a), given $I \in \Lambda_j$, by (58d) and (57)

$$\rho(I) \geq \rho(B_{\mathcal{M}}(x_I, r_0\delta^j)) \geq d_1 p_1 \delta^{jd}.$$

Hence

$$\operatorname{diam}_\rho(I) \leq 2r_1\delta^j \leq 2r_1 \left(\frac{\rho(I)}{d_1 p_1}\right)^{\frac{1}{d}},$$

so that (25a) holds true with $C_1 = 2r_1(d_1p_1)^{-\frac{1}{d}}$ and $s = 1/d$. □

Proof of Lemma 4.4.1. With the notation of the above proof. Since $E_I = 0$ if $\rho_I = 0$, we can assume that $\rho_I > 0$, so that by assumption $\mu_I > 0$ and, with (40),

$$E_I \leq \Phi_+ \frac{\rho_I}{\mu_I} \int_I \|x - c_I\|^2 \, d\mu(x) \leq \Phi_+ D_s^2 \rho_I \mu_I^{2s}.$$

Fix $q \geq 1$, as in (52)

$$\eta^{\frac{1}{q}} N_+ \leq \Phi_+^{\frac{1}{q}} D_s^{\frac{2}{q}} \sum_{I \in \Lambda_+} \rho_I^{\frac{1}{q}} \mu_I^{\frac{2s}{q}}, \tag{61}$$

and Holder inequality with $1/p + 1/q = 1$ gives

$$\begin{aligned}\sum_{I\in\Lambda_+} \rho_I^{\frac{1}{q}} \mu_I^{\frac{2s}{q}} &\leq \left(\sum_{I\in\Lambda_+} (\rho_I \mu_I^{2s})^{\frac{p}{q}}\right)^{\frac{1}{p}} \left(\sum_{I\in\Lambda_+} 1^q\right)^{\frac{1}{q}} \\ &\leq \left(\sum_{I\in\Lambda} \mu_I (\frac{\rho_I}{\mu_I})^{\frac{1}{2s+1}}\right)^{\frac{2s+1}{2s+2}} N_+^{\frac{1}{2s+2}} = \left(\sum_{I\in\Lambda} \rho_I (\frac{\rho_I}{\mu_I})^{-\frac{2s}{2s+1}}\right)^{\frac{2s+1}{2s+2}} N_+^{\frac{1}{2s+2}}\end{aligned}$$

where the last inequality follows by choosing $p/q = 1/(2s+1)$, *i.e.* $q = 2s+2$. Since the function $t \mapsto t^{-\frac{2s}{2s+1}}$ is convex on $(0, +\infty)$, Jensen inequality and the first inequality in (39) give

$$\begin{aligned}\rho_I (\frac{\rho_I}{\mu_I})^{-\frac{2s}{2s+1}} \leq \frac{\rho_I}{\mu_I} \int_I p(x)^{-\frac{2s}{2s+1}} \mathbb{1}_{|p(x)>0} \, d\mu(x) &\leq \frac{\rho_I}{\mu_I} \int_I \frac{p(x)^{\frac{1}{2s+1}}}{p(x)} d\mu(x) \leq \frac{1}{\Phi_-} \int_I p(x)^{\frac{1}{2s+1}} d\mu(x) \\ &\leq \frac{1}{\Phi_-} \int_I p(x)^{-\frac{2s}{2s+1}} \mathbb{1}_{|p(x)>0} d\mu(x).\end{aligned}$$

By replacing in (61) we get

$$\begin{aligned}N_+ &\leq \Phi_-^{-1} \left(D_s^2 \, \Phi_+ \, \|p\|_{\frac{1}{2s+1}}\right)^{\frac{1}{2s+1}} \eta^{-\frac{1}{2s+1}} \\ &\leq \Phi_-^{-1} \left(D_s^2 \, \Phi_+ \left\|p^{-2s} \mathbb{1}_{|p(x)>0}\right\|_{\frac{1}{2s+1}}\right)^{\frac{1}{2s+1}} \eta^{-\frac{1}{2s+1}}\end{aligned}$$

□

6. Statistical Tools

We recall the following probabilistic inequality based on a result of [84, 85], see also [110, Theorem 3.3.4] and [83] for concentration inequalities for Hilbert-space-valued random variables.

PROPOSITION 6.0.1. *Let $\xi_1, \ldots, \xi_n$ be a family of independent zero-mean random variables taking values in a real separable Hilbert space and satisfying*

$$\mathbb{E}[\|\xi_i\|^m] \leq \frac{1}{2} m! \Sigma^2 M^{m-2} \qquad \forall m \geq 2, \tag{62}$$

where Σ and M are two positive constants. Then, for all $n \in \mathbb{N}$ and $t > 0$

$$\mathbb{P}\left[\left\|\frac{1}{n}\sum_{i=1}^{n}\xi_i\right\| \geq t\right] \leq 2\exp\left(-\frac{nt^2}{\Sigma^2 + Mt + \Sigma\sqrt{\Sigma^2 + 2Mt}}\right) = 2\exp\left(-n\frac{\Sigma^2}{M^2} g\left(\frac{Mt}{\Sigma^2}\right)\right) \tag{63}$$

where $g(t) = \frac{t^2}{1+t+\sqrt{1+2t}}$. In particular, if ξ_i are bounded by M with probability 1, then

$$\mathbb{P}\left[\left\|\frac{1}{n}\sum_{i=1}^{n}\xi_i\right\| \geq t\right] \leq 2\exp\left(-\frac{nt^2}{4M^2}\right). \tag{64}$$

PROOF. Bound(63) is given in [84] with a wrong factor, see [85]. To show(64), note that (62) is satisfied with $\Sigma = M$. Furthermore, for all $t \leq 1$, $g(t) \geq t^2/4$, so that if $t \leq M$,

$$\mathbb{P}\left[\left\|\frac{1}{n}\sum_{i=1}^{n}\xi_i\right\| \geq t\right] \leq 2\exp\left(-\frac{nt^2}{4M^2}\right)$$

whereas, if $t > M$, (64) is trivially satisfied. □

The following concentration inequality is based on [66] and we adapt the proof of Theorem 10 in [67]. We introduce the following notation. Given a family $\xi_1, \ldots \xi_n$ of independent random variables taking value in some measurable space $\mathcal{Y}$ and a measurable positive bounded function $f : \mathcal{Y}^n \to \mathbb{R}$, for any $k = 1, \ldots, n$ set

$$\begin{aligned} V_k &= f(\xi_1, \ldots, \xi_n) - \inf_{y \in \mathcal{Y}} f(\xi_1, \ldots, \xi_{k-1}, y, \xi_{k+1}, \ldots, \xi_n) \\ &= \sup_{y \in \mathcal{Y}} \Big(f(\xi_1, \ldots, \xi_n) - f(\xi_1, \ldots, \xi_{k-1}, y, \xi_{k+1}, \ldots, \xi_n) \Big). \end{aligned}$$

PROPOSITION 6.0.2. *With the above notation, if there exist two constants $\alpha, \beta > 0$ such that*

$$\max_{k=1,\ldots,n} V_k \leq \alpha \tag{65a}$$

$$\sum_{k=1}^{n} V_k^2 \leq \beta f(\xi_1, \ldots, \xi_n) \tag{65b}$$

then, for any $t > 0$

$$\mathbb{P}\left[\left|\sqrt{f(\xi_1, \ldots, \xi_n)} - \sqrt{\mathbb{E}[f(\xi_1, \ldots, \xi_n)]}\right| > t\right] \leq 2\exp\left(-\frac{t^2}{2\max\{\alpha, \beta\}}\right). \tag{66}$$

PROOF. Set $Z_k = V_k/\alpha$ and $Z = f(\xi_1, \dots, \xi_n)/\alpha$. By construction

$$Z_k \leq 1 \qquad k = 1, \dots, n$$

$$\sum_{k=1}^{n} Z_k^2 \leq \frac{\beta}{\alpha} Z.$$

Let $\gamma = \max\{\beta/\alpha, 1\}$. Theorem 13 of [66] gives that

$$\mathbb{P}\left[\mathbb{E}\left[Z\right] - Z > t\right] \leq \exp(-\frac{t^2}{2\gamma\mathbb{E}[Z]}) \tag{67}$$

$$\mathbb{P}\left[Z - \mathbb{E}\left[Z\right] > t\right] \leq \exp(-\frac{t^2}{2\gamma\mathbb{E}[Z] + \gamma t}). \tag{68}$$

By replacing t with $2t\sqrt{\mathbb{E}[Z]}$ in (67)

$$\begin{aligned}
\exp(-\frac{2t^2}{\gamma}) &\geq \mathbb{P}\left[\mathbb{E}\left[Z\right] - 2t\sqrt{\mathbb{E}[Z]} + t^2 > Z + t^2\right] \\
&= \mathbb{P}\left[\left|\sqrt{\mathbb{E}[Z]} - t\right| > \sqrt{Z + t^2}\right] \\
&\geq \mathbb{P}\left[\sqrt{\mathbb{E}[Z]} - \sqrt{Z} > 2t\right]
\end{aligned}$$

since

$$\sqrt{\mathbb{E}[Z]} - t \leq \left|\sqrt{\mathbb{E}[Z]} - t\right| \leq \sqrt{Z + t^2} \leq \sqrt{Z} + t$$

provided that $\left|\sqrt{\mathbb{E}[Z]} - t\right| \leq \sqrt{Z + t^2}$. Hence

$$\mathbb{P}\left[\sqrt{\mathbb{E}[Z]} - \sqrt{Z} > t\right] \leq \exp(-\frac{t^2}{2\gamma}). \tag{69}$$

Setting $\frac{t^2}{2\mathbb{E}[Z]+t} = 2\tau^2$, bound (68) gives

$$\begin{aligned}
\exp(-\frac{2\tau^2}{\gamma}) &\geq \mathbb{P}\left[Z - \mathbb{E}[Z] > \tau^2 + \sqrt{\tau^4 + 4\tau^2\mathbb{E}[Z]}\right] \\
&\geq \mathbb{P}\left[Z - \mathbb{E}[Z] > 4\tau^2 + 4\tau\sqrt{\mathbb{E}[Z]}\right] \\
&= \mathbb{P}\left[Z > \left(\sqrt{\mathbb{E}[Z]} + 2\tau\right)^2\right] \\
&= \mathbb{P}\left[\sqrt{Z} - \sqrt{\mathbb{E}[Z]} > 2\tau\right],
\end{aligned}$$

so that, setting $\tau = t/2$,

$$\mathbb{P}\left[\sqrt{Z} - \sqrt{\mathbb{E}[Z]} > t\right] \leq \exp(-\frac{t^2}{2\gamma}). \tag{70}$$

Bounds (69) and (70) imply that

$$\mathbb{P}\left[\left|\sqrt{Z} - \sqrt{\mathbb{E}[Z]}\right| > t\right] \leq 2\exp(-\frac{t^2}{2\gamma}),$$

and, by replacing t with $t/\sqrt{\alpha}$,

$$\mathbb{P}\left[\left|\sqrt{f(\xi_1, \dots, \xi_n)} - \sqrt{\mathbb{E}[f(\xi_1, \dots, \xi_n)]}\right| > t\right] \leq 2\exp(-\frac{t^2}{2\alpha\gamma}).$$

where $\alpha\gamma = \alpha \max\{\beta/\alpha, 1\} = \max\{\beta, \alpha\}$. □

7. Proofs of Stability Results

The proof of Theorem 3.5.1 is based on the estimate of Proposition (55a), plus borrows some ideas from [12, 59] adapted to our setting through some intermediate results. For sake of clarity let $\hat{P}_\eta = \widehat{P}_{\hat{\Lambda}_\eta}$.

Proof of Thm. 3.5.1. Consider the following decomposition

$$x - \widehat{P}_{\hat{\Lambda}_\eta}(x) = \left(x - P_{\Lambda(\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta})}(x)\right) + \left(P_{\Lambda(\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta})}(x) - P_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2})}(x)\right) + \\ + \left(P_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2})}(x) - \widehat{P}_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{2\eta})}(x)\right) + \left(\widehat{P}_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{2\eta})}(x) - \widehat{P}_{\Lambda(\hat{\mathcal{T}}_\eta)}(x)\right),$$

which holds for all $x \in \mathcal{X}$. Since

$$\left\| \sum_{i=1}^4 v_i \right\|^2 \leq 4 \sum_{i=1}^4 \|v_i\|^2 \qquad v_1, \dots, v_4 \in \mathbb{R}^D,$$

it holds that

$$\mathcal{E}[\widehat{P}_{\hat{\Lambda}_\eta}] \lesssim \underbrace{\mathcal{E}[P_{\Lambda(\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta})}]}_{\text{A}} + \underbrace{\int_{\mathcal{X}} \left\| P_{\Lambda(\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta})}(x) - P_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2})}(x) \right\|^2 d\rho(x)}_{\text{B}} \\ + \underbrace{\int_{\mathcal{X}} \left\| P_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2})}(x) - \widehat{P}_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2})}(x) \right\|^2 d\rho(x)}_{\text{C}} \\ + \underbrace{\int_{\mathcal{X}} \left\| \widehat{P}_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{2\eta})}(x) - \widehat{P}_{\Lambda(\hat{\mathcal{T}}_\eta)}(x) \right\|^2 d\rho(x)}_{\text{D}}.$$

We bound the four terms.

A) Since $\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta} \supset \mathcal{T}_{2\eta}$, $\Lambda(\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta})$ is a partition finer than $\Lambda_{2\eta}$, then

$$\mathcal{E}[P_{\Lambda(\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta})}] \leq \mathcal{E}[P_{\Lambda_{2\eta}}] \lesssim \eta^{\frac{4\sigma}{2\sigma+1}},$$

where the last inequality is a consequence of (34).

B) Bound (81a) implies that the term B is zero with probability greater than $1 - p_B$ where

$$p_B \lesssim (n^\gamma + \eta^{-\frac{2}{2\sigma+1}}) \exp(-c_a n \eta^2).$$

C) Since $\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2}) \subset \mathcal{T}_{\eta/2} \cup \Lambda_{\eta/2} = \mathcal{I}$ and $\sharp\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2}) \leq \sharp\Lambda_{\eta/2} = N$, by (72) term C is bounded by $t^* = \eta^{\frac{4\sigma}{2\sigma+1}} t$ with probability greater than $1 - p_C$ with

$$p_C = 2\sharp\mathcal{I} \exp\left(-\frac{nt^*}{4N}\right) \lesssim \eta^{-\frac{2}{2\sigma+1}} \exp\left(-c_{\mathbb{T}} n \eta^2 t\right) \tag{71}$$

where the second inequality is a consequence of (55a) and (55b), and $c_{\mathbb{T}} > 0$ is a suitable constant depending on the partition tree $\mathbb{T}$.

D) By (81b) term D is zero with probability greater that $1 - p_D$ where

$$p_D \lesssim \eta^{-\frac{2}{2\sigma+1}} \exp(-c_a n \eta^2).$$

If follows that with probability greater than $1-(p_B+p_C+p_B)$

$$\mathcal{E}[\widehat{P}_{\widehat{\Lambda}_\eta}] \lesssim \underbrace{\eta^{\frac{4\sigma}{2\sigma+1}}}_{A} + \underbrace{\eta^{\frac{4\sigma}{2\sigma+1}} t}_{C}$$

i.e.

$$\mathbb{P}\left[\mathcal{E}[\widehat{P}_{\widehat{\Lambda}_\eta}] \gtrsim \eta^{\frac{4\sigma}{2\sigma+1}}(1+t)\right] \lesssim \underbrace{(n^\gamma + \eta^{-\frac{2}{2\sigma+1}}) \exp\left(-c_a n\eta^2 t\right)}_{p_A+p_D} + \underbrace{\eta^{-\frac{2}{2\sigma+1}} \exp\left(-c_{\mathbb{T}} n\eta^2 t\right)}_{p_C}$$

which gives (29). □

Proof of Cor. 3.5.2. Since $\eta_n^2 = \frac{(\gamma+\beta)\ln n}{c_a n}$, then bound (29) gives (31a) since

$$(n^\gamma + \left(\frac{c_a n}{(\gamma+\beta)\ln n}\right)^{\frac{1}{2\sigma+1}}) \exp\left(-c_a n\eta_n^2\right) \lesssim n^\gamma n^{-(\gamma+\beta)} = n^{-\beta}$$

$$\left(\frac{c_a n}{(\gamma+\beta)\ln n}\right)^{\frac{1}{2\sigma+1}} \exp\left(-c_{\mathbb{T}} n\eta_n^2 t\right) \lesssim n n^{-\bar{c}_{\mathbb{T}} t} = n^{1-\bar{c}_{\mathbb{T}} t}$$

where $\bar{c}_{\mathbb{T}} = c_{\mathbb{T}} c_a^{-1}(\gamma+\beta)$. Eq. (31b) is clear. □

Proof of Prop. 3.6.1. Given $I \in \mathbb{T}$, by (50) and (25b),

$$\sum_{J \in \mathcal{C}^{N+1}(I)} \mathcal{E}_J \lesssim a\, b^{-2(j_I+N+1)}$$

so that

$$\lim_{N\to+\infty} \sum_{J\in\mathcal{C}^{N+1}(I)} \mathcal{E}_J = 0$$

and, by taking the limit in (48),

$$\mathcal{E}_I = \sum_{k=0}^{+\infty} \sum_{J\in\mathcal{C}^k(I)} \epsilon_J^2.$$

Set $\mathcal{T}_k = \mathcal{T}_{\eta/2^k}$ for all $k \in \mathbb{N}$, then

$$\begin{aligned}
\mathcal{E}(P_{\Lambda_\eta}) &= \sum_{I\in\Lambda_\eta} \mathcal{E}_I = \sum_{I\in\Lambda_\eta} \sum_{k=0}^{+\infty} \sum_{J\in\mathcal{C}^k(I)} \epsilon_J^2 = \sum_{J\notin\mathcal{T}_\eta} \epsilon_J^2 \\
&= \sum_{k=0}^{+\infty} \sum_{J\in\mathcal{T}_{k+1}\setminus\mathcal{T}_k} \epsilon_J^2 \le \sum_{k=0}^{+\infty} \sharp\mathcal{T}_{k+1} (\frac{\eta}{2^k})^2 \\
&\lesssim \sum_{k=0}^{+\infty} \frac{\eta^2}{2^{2k}} (\frac{\eta}{2^{k+1}})^{-\frac{2}{2\sigma+1}} = \eta^{\frac{4\sigma}{2\sigma+1}} \sum_{k=0}^{+\infty} 4^{\frac{k+1}{2\sigma+1}-k} \lesssim \eta^{\frac{4\sigma}{2\sigma+1}}
\end{aligned}$$

where the first inequality is a consequence of the fact that $\epsilon_J < (\frac{\eta}{2^k})^2$ if $J \notin \mathcal{T}_k$, the second inequality follows from (55a), whereas the last inequality holds since the series $\sum_{k=0}^{+\infty} 4^{-\frac{2\sigma k-1}{2\sigma+1}}$ converges. □

7.1. C-term: sample error. The following result bounds the sample error for a given partition.

PROPOSITION 7.1.1. *Fix a data-independent subset $\mathcal{I} \subset \mathbb{T}$ of cells and $N > 0$. Given a partition $\widehat{\Lambda} \subset \mathcal{I}$ (possibly depending on the data) such that $\sharp\widehat{\Lambda} \leq N$, for any $t > 0$,*

$$\mathbb{P}\left[\int_{\mathcal{X}} \left\| P_{\widehat{\Lambda}}(x) - \widehat{P}_{\widehat{\Lambda}}(x) \right\|^2 d\rho(x) > t\right] \leq 2\,\sharp\mathcal{I} \exp\left(-\frac{nt}{8N}\right) \tag{72}$$

PROOF. Consider the following event

$$\Omega = \bigcup_{I \in \mathcal{I}} \{\sqrt{\rho_I}\, \|\widehat{c}_I - c_I\| > t\},$$

which is well-defined since $\mathcal{I}$ does not depend on the data $X_1, \dots, X_n$. By union bound

$$\mathbb{P}[\Omega] \leq \sharp\mathcal{I} \sup_{\substack{I \in \mathcal{I} \\ \rho_I > 0}} \mathbb{P}[\|\widehat{c}_I - c_I\| > \frac{t}{\sqrt{\rho_I}}]. \tag{73}$$

Fix $I \in \mathcal{I}$ with $\rho_I > 0$. The tower property with respect to the binomial random variable n_I gives

$$\mathbb{P}[\|\widehat{c}_I - c_I\| > t] = \sum_{k=0}^{n} \binom{n}{k} \rho_I^k (1-\rho_I)^{n-k}\, \mathbb{P}[\|\widehat{c}_I - c_I\| > t \mid n_I = k].$$

Conditionally to the event $\{n_I = k\}$ with $k > 0$, up to a permutation of the indexes, we can assume that $X_1, \dots, X_k \in I$ and $X_{k+1}, \dots, X_n \notin I$. Furthermore,

$$\widehat{c}_I - c_I = \frac{1}{k}\sum_{i=1}^{k}(X_i - c_I) = \frac{1}{k}\sum_{i=1}^{k}\xi_i$$

where $\xi_1, \dots, \xi_k$ are independent zero mean random vectors bounded by $M = \operatorname{diam}_\rho(I)$ almost surely by (45a). Hence, by (64)

$$\mathbb{P}[\|\widehat{c}_I - c_I\| > t \mid n_I = k] \leq 2\exp(-\frac{kt^2}{4\operatorname{diam}_\rho(I)^2}),$$

which holds true also if $k = 0$. Hence,

$$\begin{aligned}
\mathbb{P}[\|\widehat{c}_I - c_I\| > t] &\leq 2\sum_{k=0}^{n}\binom{n}{k}(\rho_I \exp(-\frac{t^2}{4\operatorname{diam}_\rho(I)^2}))^k(1-\rho_I)^{n-k} \\
&= 2\left(1 - \rho_I\left(1 - \exp(-\frac{t^2}{4\operatorname{diam}_\rho(I)^2})\right)\right)^n \\
&\leq 2\exp\left(-n\rho_I\left(1 - \exp(-\frac{t^2}{4\operatorname{diam}_\rho(I)^2})\right)\right)
\end{aligned}$$

where in the fourth term we use the bound $(1-\tau)^n \leq \exp(-n\tau)$ with $0 \leq \tau \leq 1$. Since

$$1 - \exp(-\tau) \geq \frac{\tau}{2} \qquad \text{for all } \tau \leq 1,$$

then, for all $t \leq \operatorname{diam}_\rho(I)$,

$$\mathbb{P}[\|\widehat{c}_I - c_I\| > t] \leq 2\exp\left(-n\rho_I\frac{t^2}{8\operatorname{diam}_\rho(I)^2}\right) \leq 2\exp\left(-n\rho_I\frac{t^2}{8}\right). \tag{74}$$

If $\operatorname{diam}_\rho(I) < t \le \operatorname{diam}\mathcal{X} \le 1$, by (45c), clearly $\mathbb{P}[\|\widehat{c}_I - c_I\| > t \mid n_I > 0] = 0$, so that

$$\mathbb{P}[\|\widehat{c}_I - c_I\| > t] = \mathbb{P}[\|\widehat{x}_I^* - c_I\| > t] \mid n_I = 0]\mathbb{P}\left[n_I = 0\right] \le (1-\rho_I)^n \le \exp(-n\rho_I) \le 2\exp\left(-n\rho_I\frac{t^2}{8}\right),$$

where the last bound holds true for any $t \le 2\sqrt{2}$. Finally, if $t > \operatorname{diam}\mathcal{X}$, as above

$$\mathbb{P}[\|\widehat{c}_I - c_I\| > t] = \mathbb{P}[\|\widehat{x}_I^* - c_I\| > t] \mid n_I = 0]\mathbb{P}\left[n_I = 0\right] = 0$$

since $\widehat{x}_I^*, c_I \in \mathcal{X}$, compare with (45c). It follows that (74) holds true for all $t > 0$. From (73)

$$\mathbb{P}[\Omega] \le 2\,\sharp\mathcal{I}\exp\left(-\frac{nt^2}{8}\right).$$

Since $\widehat{\Lambda} \subset \mathcal{I}$, on the complement of Ω,

$$\int_{\mathcal{X}} \left\|P_{\widehat{\Lambda}}(x) - \widehat{P}_{\widehat{\Lambda}}(x)\right\|^2 d\rho(x) = \sum_{\substack{I\in\widehat{\Lambda}\\ \rho_I>0}} \rho_I \left\|\widehat{c}_I - c_I\right\|^2 \le Nt^2,$$

and bound (72) follows by replacing t with $\sqrt{t/N}$. □

Remark 7.1.2. *By inspecting the proof, it turns out that the assumption $\hat{x_I}^* \in \mathcal{X}$ is only needed in this proposition and it can be replaced by the condition that $\inf_{x\in\mathcal{X}} \|\hat{x}_I^* - x\| \le 1$, so that $\|\hat{x}_I^* - c_I\| \le 2$.*

7.2. B and D terms: Stability of P_Λ with respect to the partition. The following result is well-known.

Lemma 7.2.1. *Given a cell $I \in \mathbb{T}$ with $\rho_I > 0$, for all $t > 0$*

$$\mathbb{P}\left[\left|\frac{n_I}{n} - \rho_I\right| \ge \rho_I t\right] \le 2\exp\left(-\frac{n\rho_I t^2}{2(1+t/3)}\right) \le 2\exp\left(-\frac{n\rho_I t^2}{M_I}\right) \tag{75}$$

where $M_I = 2/3\max\{4, (1+2\rho_I)/\rho_1\}$.

Proof. We apply the Bernstein inequality, see [17, Corollary 2.11], to the family of independent random variables $\mathbb{1}_I(X_1), \dots, \mathbb{1}_I(X_n)$, which satisfy

$$\begin{aligned}
&\mathbb{E}[\mathbb{1}_I(X_i)] = \rho_I && i = 1, \dots, n\\
&\sum_{i=1}^n \mathbb{E}[\mathbb{1}_I(X_i)^2] = n\rho_I\\
&\sum_{i=1}^n \mathbb{E}[\mathbb{1}_I(X_i)^m] = n\rho_I \le \frac{n\rho_I}{2} m!(\frac{1}{3})^{m-2} && m \in \mathbb{N},\ m \ge 3,
\end{aligned}$$

then,

$$\mathbb{P}\left[|n_I - n\rho_I| \ge n\rho_I t\right] \le 2\exp\left(-\frac{(n\rho_I t)^2}{2(n\rho_I + n\rho_I t/3)}\right) = 2\exp\left(-\frac{n\rho_I t^2}{2(1+t/3)}\right).$$

Observing that

$$|n_I - n\rho_I| \le n\max\{\rho_I, 1-\rho_I\}.$$

If $t > \max\{1, 1/\rho_I - 1\} = t^*$, then $\mathbb{P}\left[|n_I - n\rho_I| \ge n\rho_I t\right] = 0$. If $t \le \max\{1, 1/\rho_I - 1\} = t^*$, it holds that

$$2(1+t/3) \le 2(1+t^*/3) = M_I,$$

so that the second bound in (75) is clear. □

The following lemma provides a concentration inequality of $\sqrt{\frac{n_I}{n}}$.

LEMMA 7.2.2. *Given a cell $I \in \mathbb{T}$ with $\rho_I > 0$, for all $t > 0$*

$$\mathbb{P}\left[\left|\sqrt{\frac{n_I}{n}} - \sqrt{\rho_I}\right| \geq t\right] \leq 2\exp\left(-\frac{nt^2}{2}\right). \tag{76}$$

PROOF. We apply Proposition 6.0.2 with $\mathcal{Y} = \{0, 1\}$

$$\xi_i = \mathbb{1}_I(X_i) \qquad f(y_1, \ldots, y_n) = \frac{1}{n}\sum_{i=1}^{n} y_i,$$

where f is clearly bounded, and

$$f(\xi_1, \ldots, \xi_n) = \frac{n_I}{n} \qquad \mathbb{E}\left[f(\xi_1, \ldots, \xi_n)\right] = \rho_I.$$

Given $k = 1, \ldots, n$, it holds that

$$V_k(\xi_1, \ldots, \xi_n) = \frac{1}{n}\sup_{y\in\mathcal{Y}}\left(\mathbb{1}_I(X_i) - y\right) = \frac{1}{n}\mathbb{1}_I(X_i) \leq \frac{1}{n},$$

so that $\alpha = 1/n$. Furthermore

$$\sum_k V_k^2(\xi_1, \ldots, \xi_n) = \frac{1}{n^2}\sum_k \mathbb{1}_I(X_i) = \frac{1}{n} f(\xi_1, \ldots, \xi_n)$$

then $\beta = 1/n$ and Eq. (66) implies (76). □

The following lemma shows that, given a cell $I \in \mathbb{T}$, $\widehat{\epsilon}_I$ concentrates around ϵ_I.

LEMMA 7.2.3. *Given $I \in \mathbb{T}$, for all $t > 0$*

$$\mathbb{P}\left[|\widehat{\epsilon}_I - \epsilon_I| > t\right] \leq 2\ell\exp\left(-n\frac{t^2}{64\ell\,\mathrm{diam}_\rho(I)^2}\right). \tag{77}$$

where $\ell = 1 + \sharp\mathcal{C}(I)$.

PROOF. Fix $I \in \mathbb{T}$. If $\rho_J = 0$ for some $J \in \mathcal{C}(I)$, then almost surely $X_i \notin J$ for all $i = 1, \ldots, n$ and, hence, $n_J = 0$, so that both $\widehat{\epsilon}_I$ and ϵ_I do not depend on the children J. Hence, without loss of generality, we can assume that $\rho_J > 0$ for all $J \in \mathcal{C}(I)$.

Let $\ell = \sharp\mathcal{C}(I) + 1$. Set $L^2(\mathcal{C}(I)) = \mathbb{R}^{\ell-1}$, regarded as Euclidean vector space whose norm is denoted by $\|v\|_2$. Define $v, \widehat{v}, \widehat{w} \in L^2(\mathcal{C}(I))$

$$\begin{aligned} v(J) &= \sqrt{\rho_J}\,\|c_J - c_I\| \\ \widehat{v}(J) &= \sqrt{\frac{n_J}{n}}\,\|\widehat{c}_J - \widehat{c}_I\| \\ \widehat{w}(J) &= \sqrt{\rho_J}\,\|\widehat{c}_J - \widehat{c}_I\|\,. \end{aligned}$$

Then

$$|\widehat{\epsilon}_I - \epsilon_I| = |\|\widehat{v}\|_2 - \|v\|_2| \leq \|\widehat{v} - v\|_2 \leq \|\widehat{w} - v\|_2 + \|\widehat{v} - \widehat{w}\|_2\,. \tag{78}$$

We now bound the first term. Set $\mathcal{I} = \mathcal{C}(I) + \{I\}$ and $\sharp\mathcal{I} = \ell$.

$$\begin{aligned}\|\widehat{w} - v\|_2^2 &= \sum_{J\in\mathcal{C}(I)} \rho_J \left|\|\widehat{c}_J - \widehat{c}_I\| - \|c_J - c_I\|\right|^2 \\ &\le 2\sum_{J\in\mathcal{C}(I)} \rho_J \left|\|\widehat{c}_J - c_J\| + \|\widehat{c}_I - c_I\|\right|^2 \\ &\le 2\sum_{J\in\mathcal{I}} \rho_J \|\widehat{c}_J - c_J\|^2 \\ &\le 2\ell \max_{J\in\mathcal{I}} \rho_J \|\widehat{c}_J - c_J\|^2,\end{aligned}$$

Setting

$$\Omega_I = \bigcup_{J\in\mathcal{I}} \left\{\|\widehat{c}_J - c_J\| > \frac{t}{\sqrt{2\ell\rho_J}}\right\},$$

bound (74) gives

$$\mathbb{P}\left[\Omega_I\right] \le 2\ell \exp\left(-\frac{nt^2}{16\ell \operatorname{diam}_\rho(I)^2}\right),$$

so that

$$\mathbb{P}\left[\|\widehat{w} - v\|_2 > t\right] \le 2\ell \exp\left(-\frac{nt^2}{16\ell \operatorname{diam}_\rho(I)^2}\right). \tag{79}$$

We now bound the second term in (78). Set

$$\Omega'_I = \bigcup_{J\in\mathcal{C}(I)} \left\{\left|\sqrt{\frac{n_J}{n}} - \sqrt{\rho_J}\right| > \frac{t}{\sqrt{\ell}\operatorname{diam}_\rho(I)}\right\},$$

bound (75) gives

$$\mathbb{P}\left[\Omega'_I\right] \le 2\ell \exp\left(-\frac{nt^2}{2\ell \operatorname{diam}_\rho(I)^2}\right),$$

On the complement on Ω'_I

$$\begin{aligned}\|\widehat{w} - \widehat{v}\|_2^2 &= \sum_{J\in\mathcal{C}(I)} \|\widehat{c}_J - \widehat{c}_I\|^2 \left(\sqrt{\frac{n_J}{n}} - \sqrt{\rho_J}\right)^2 \\ &\le \ell \operatorname{diam}_\rho(I)^2 \sup_{J\in\mathcal{C}(I)} \left|\sqrt{\frac{n_J}{n}} - \sqrt{\rho_J}\right|^2 \le t^2.\end{aligned}$$

Hence

$$\mathbb{P}\left[\|\widehat{w} - \widehat{v}\|_2 > t\right] \le 2\ell \exp\left(-\frac{nt^2}{2\ell \operatorname{diam}_\rho(I)^2}\right). \tag{80}$$

Inequality (78) with bounds (79) and (80) implies (77). □

The next proposition shows that P_Λ is stable under suitable small perturbations of the partition Λ.

PROPOSITION 7.2.4. *For any* $\eta > 0$

$$\mathbb{P}\left[\left(\int_{\mathcal{X}} \left\|P_{\Lambda(\hat{\mathcal{T}}_\eta \cup \mathcal{T}_{2\eta})}(x) - P_{\Lambda(\hat{\mathcal{T}}_\eta \cap \mathcal{T}_{\eta/2})}(x)\right\|^2 d\rho(x)\right) > 0\right] \lesssim (n^\gamma + \eta^{-\frac{2}{2\sigma+1}}) \exp(-c_a n\eta^2) \tag{81a}$$

and

$$\mathbb{P}\left[\left(\int_{\mathcal{X}}\left\|P_{\Lambda(\hat{\mathcal{T}}_\eta\cup\mathcal{T}_{2\eta})}(x)-P_{\Lambda(\hat{\mathcal{T}}_\eta)}(x)\right\|^2 d\rho(x)\right)>0\right]\lesssim \eta^{-\frac{2}{2\sigma+1}}\exp(-c_a n\eta^2). \tag{81b}$$

where

$$c_a=\frac{1}{128(a+1)}. \tag{81c}$$

Proof. Recalling the definition of j_n given by (21), set $\mathcal{T}_n^*=\bigcup_{j\le j_n}\Lambda_j$, which is a subtree with

$$\sharp\mathcal{T}_n^*\le\sum_{j=0}^{j_n}a^j=\frac{a^{j_n+1}-1}{a-1}\lesssim n^\gamma, \tag{82}$$

and, by construction, $\widehat{T_\eta}\subset\mathcal{T}_n^*$. The probability of the event in the left hand side of (81a) is clearly bounded by the probability of the event

$$\{\hat{\mathcal{T}}_\eta\cap\mathcal{T}_{\eta/2}\subsetneq\hat{\mathcal{T}}_\eta\cup\mathcal{T}_{2\eta}\}=\bigcup_{I\in\mathbb{T}}\{I\in\hat{\mathcal{T}}_\eta\wedge I\notin\mathcal{T}_{\eta/2}\}\cup\{I\in\mathcal{T}_{2\eta}\wedge I\notin\hat{\mathcal{T}}_\eta\}.$$

About the first term, we observe that if $I\in\hat{\mathcal{T}}_\eta\subset\mathcal{T}_n^*$, then there exist $k\ge0$ and $J\in\mathcal{C}^k(I)\cap\mathcal{T}_n^*$ such that $\widehat{\epsilon}_J\ge\eta$ and, since $I\notin\mathcal{T}_{\eta/2}$ and $J\in\mathcal{C}^k(I)$, then $\epsilon_J<\frac{\eta}{2}$, so that

$$\bigcup_{I\in\mathcal{T}_n^*}\{I\in\hat{\mathcal{T}}_\eta\wedge I\notin\mathcal{T}_{\eta/2}\}\subset\bigcup_{J\in\mathcal{T}_n^*}\{\widehat{\epsilon}_J\ge\eta\wedge\epsilon_J<\frac{\eta}{2}\}\subset\bigcup_{J\in\mathcal{T}_n^*}\{|\widehat{\epsilon}_J-\epsilon_J|>\frac{\eta}{2}\}.$$

By union bound and (77) with $t=\eta/2$, $\mathrm{diam}_\rho(I)\le1$ and $\ell\le a+1$ give

$$\begin{aligned}\mathbb{P}\left[\bigcup_{I\in\mathcal{T}_n^*}\{\widehat{\epsilon}_I\ge\eta\wedge\epsilon_I<\frac{\eta}{2}\}\right]&\lesssim\sharp\mathcal{T}_n^*\exp\left(-c_a n\eta^2\right)\\&\lesssim n^\gamma\exp(-c_a n\eta^2)\end{aligned} \tag{83}$$

where the second inequality is a consequence of (82) and c_a is given by (81c).

By a similar argument

$$\{I\in\mathcal{T}_{2\eta}\wedge I\notin\hat{\mathcal{T}}_\eta\}\subset\bigcup_{J\in\mathcal{T}_{2\eta}}\{|\widehat{\epsilon}_J-\epsilon_J|>\eta\}.$$

By union bound and (77) with $t=\eta$ and $\mathrm{diam}_\rho(I)\le1$ give

$$\mathbb{P}\left[\bigcup_{I\in\mathcal{T}_{2\eta}}\{\epsilon_I\ge2\eta\wedge\widehat{\epsilon}_I<\eta\}\right]\lesssim\sharp\mathcal{T}_{2\eta}\exp\left(-c_a n\eta^2\right)\lesssim\eta^{-\frac{2}{2\sigma+1}}\exp(-c_a n\eta^2), \tag{84}$$

where the second inequality is a consequence of (55a).

By (84) and (83), we get (81a). The proof of (81b) can be deduced reasoning as for (84). □

8. Possible Improvements

In this section we review a few points in which the proofs seem to suggest that a slightly finer analysis could reached. These improvements have not been pursued, either because we weren't able to work out them in full detail, or because the main result is already competitive as it is, and further complications would only make it less readable, or because of time constraints.

8.1. Logarithmic factor in Proposition (5.1.9).

The cardinality of the subtree $\mathcal{T}_\eta \subset \mathbb{T}$ is evaluated by a rough estimate which considers the tree as it was made of a number of disjoint branches root-to-leaf, all of the same length and ignoring their overlaps. It is clear that the logarithmic factor in the result stems from this procedure, while an analysis taking into account the overlaps could potentially drop the logarithmic factor. At the moment it is not clear how exactly to perform such an analysis. Since this is not the only logarithmic factor in our main result, this change by itself would not produce a substantial improvement.

8.2. Bounding $\mathrm{diam}_\rho(I)$ in Proposition (7.1.1).

The decay of $\mathrm{diam}_\rho(I)$ with respect to j_I, as well as with respect to ρ_I, is determined due to Assumption (3.3.1). Proposition (7.1.1) uses the rough estimate:

$$\frac{\rho_I}{\mathrm{diam}_\rho^2(I)} \geq \rho_I,$$

which does not exploit the decay of $\mathrm{diam}_\rho(I)$. Since the assumptions contain constants on whose value there is no control, it is difficult to take advantage of the mentioned decay.

Again on Proposition (7.1.1), the choice to have $\widehat{x}_I^*$ possibly falling outside of I has only be contemplated because it appears in [59], but does not make sense algorithmically, since the outcome can only perform worse, and choosing $\widehat{x}_I^* \in I$ doesn't seem to entail much computational complexity.

CHAPTER 2

Matroidal Structures in Graph Signal-Processing

1. Problem Statement

Graph Signal Processing is a teeming scientific avenue, where a problem of practical interest meets theoretical challenges, that pertain to branches of Mathematics as far apart as Analysis and Combinatorics, with the further appeal of being relatively recent. A space, as simple as the finite-dimensional vector space of signals on a finite set of vertexes, is decorated with the relations between the vertexes given by the weighted adjacency edges of a graph. These naturally give raise to a linear operator that to many regards resembles a Fourier transform, mirroring the graph domain into a frequency domain. On the one hand this structure asymptotically tends to mimic analysis on locally compact groups or manifolds, but on the other hand its discrete nature triggers a whole new scenario of algebraic phenomena, that don't quite have any obvious continuous counterpart. Hints towards making sense of the laid out framework are objects that already embody a discrete nature in continuous setting, such as measures with discrete support in time and frequency, also called *Dirac combs*. While these measures are key towards formulating sampling theorems and constructing wavelet frames in time-frequency Analysis, in the graph-frequency setting these boil down to distinguished combinatorial objects, the so called *Circuits* of a matroid, corresponding to the Fourier transform operator. In this study I strove to throw as many bridges as I could between the Combinatorics of this matroid and the properties of graph signals; as expressed, for instance, by Propositions 3.5.2, 3.5.4 and 3.5.5. This endeavour is, to the best of my knowledge, unprecedented. In a particularly symmetric case, corresponding to Cayley graphs of finite abelian groups, the Dirac combs are proven, through Theorems 5.2.7 and 5.2.11 and Corollary 5.2.12, to completely describe the so called *lattice of cyclic flats*, exhibiting in this case the strong property of being atomistic, among other properties. This is a strikingly concise description of the matroid, that opens many questions concerning how this highly regular structure relaxes into more general instances. Lastly, a related problem concerning the combinatorial interplay between a Fourier operator and its Spectrum is described, provided with some glimmers, namely Proposition 6.1.2 and Remark 6.2.3, towards its future development.

2. Background

In this section I first review the recent literature about Graph Signal-Processing, focusing on the structural aspects that motivated my attempts. In particular three fundamental branches of the research about Graph Signal-Processing will be presented, namely Sampling theory, Uncertainty Principles, and Bank Filters design; all of three turn out to potentially benefit from the framework of this thesis. This sections ends with a review of the basic ideas underlying the rest of this thesis.

2.1. Graph Signal-Processing.

The term Graph Signal refers to a vector indexed by the vertices $V = \{1, ..., n\}$ of a weighted undirected graph $G(V, A)$; $A = \{A_{ij}\}_{i,j\in V}$ is a real symmetric matrix called *adjacency matrix*, whose non-zero entry A_{ij} defines a weighted edge between vertices i, j. Let the matrix $\mathcal{L} = D - A$ be called Graph Laplacian, where $D = \mathrm{diag}(\{\sum_j A_{ij}\}_i)$; let $\mathcal{V}$ be the vectorial space of graph signals. Examples are ubiquitous in the applications. The vertexes of a graph may represent any type of network, for example people in a community, brain regions in neuronal networks, or stations in transportation networks. Data on a graph may be scalars defined on each nodes, which form a graph signal [74]. Extending classical signal processing methods to graph signals is the purpose of the emerging field of Graph Signal Processing [97, 96, 89]. The cornerstone of Graph Signal Processing is using $\mathcal{L}$ as a surrogate for the classical Laplace operator on euclidean spaces. This choice is usually justified by analogy with the situation of the path-graph, considered as regular discretisation of the real line $\mathbb{R}$, in which $\mathcal{L}$ exactly matches to the stencil approximation of the second derivative. Indeed, on $\mathbb{R}^N$, using the notation $H = \{\pm h\underline{e}_j\}_{j=1,...,N}$, $\underline{x} \in \mathbb{R}^N$ one has

$$\Delta f(\underline{x}) = \sum_{j=1}^{N} \frac{\partial^2}{\partial x_j^2} f(\underline{x}) = -\lim_{h\to 0} \frac{1}{h^2}\Big(2Nf(\underline{x}) - \sum_{\underline{h}\in H} f(\underline{x} + \underline{h})\Big)$$

Given a discrete abelian group G and a symmetric subset $H \subset G, H = -H$ that generates G, the corresponding Cayley Graph has vertexes $V = G$ and adjacency matrix:

$$A_{ij} = \begin{cases} 1 & \text{if } j - i \in H \\ 0 & \text{otherwise.} \end{cases}$$

The graph Laplacian in this case takes the form:

$$\mathcal{L}f(x) = |H|f(x) - \sum_{h\in H} f(x + h) = \big(|H| - A\big)f(x),$$

so that the analogy is strict. More generally, there is a whole literature addressing convergence issues of $\mathcal{L}$ to the Laplace-Beltrami operator for graphs defined as discretisations of manifolds [105]. However many graphs of interest are by no means discretizations of manifolds.

The choice of $\mathcal{L}$ allows to define a Graph Fourier Transform U^*. Let $\{u_k \in \mathcal{V}\}_{k\in W}$ be an orthonormal basis of eigenvectors of $\mathcal{L}$, having ordered eigenvalues $\{\lambda_k\}_{k\in W}$, where $W = \{1, ..., n\}$ is referred to as *frequency domain*, being the eigenvalues often called *frequencies* (whereas V is called *graph domain*). Then U will be the unitary matrix having $\{u_k\}_{k\in W}$ as rows, and U^* its adjoint.

Many processing algorithms designed for graph signals attempt to leverage intuition from time-frequency analysis by generalizing fundamental operators and transforms to the graph setting.

While there is still limited theory, a solid amount of empirical research shows that the framework is beneficial in various applications.

Among the intuitions at the core of the time-frequency analysis folklore, some do extend naturally to the *graph-frequency* setting, while others do not. Unlike the classical Fourier modes, which have global support, some graph Fourier modes have the surprising potential of being localized on very few nodes. The farther a graph is far from a regular grid, the higher is the chance to have a few localized Fourier modes, see for example [97, 1, 71].

A remarkable mismatch is the lack of a shift-invariant notion of translation for graph signals.

2.2. Sampling.

Within Signal-Processing, a cornerstone towards conveniently handling signals is provided by Sampling techniques, that consist in measuring a signal on a reduced subset of the domain, carefully chosen to enable stable reconstructions. Classically, sampling a continuous signal $x(t)$ consists in measuring a countable sequence of its values, $\{x(t_j)\}_{j\in\mathcal{Z}}$, that ensures recovery of the signal $x(t)$, under a given smoothness model [107]. Smoothness assumptions are often defined in terms of the signal's Fourier transform. For example, Shannon's famous sampling theorem [94] is classically formulated for regular sampling on $\frac{1}{\omega_s}\mathcal{Z} \subset \mathbb{R}$. This formulation requires the use of a well defined translation operator, which is in general not available for graph signals.

Despite that, the sampling theorem still admits a natural instance in the graph signal processing setting, by simple linear algebra arguments; these provide in fact a more general formulation than one might expect. On the one hand, in the continuous time-frequency analysis, a signal $f \in L_2(\mathbb{R})$ is ω_s-band-limited whenever its Fourier transform $\mathscr{F}\{f\}$ is supported on an interval $[-\omega_s/2, \omega_s/2]$ of the frequency domain. This condition constrains the smoothness of f by preventing high-frequency components to be allowed. On the other hand, as mentioned above, for graph signals the smoothness intuition doesn't necessarily carry over, and it is natural to consider a more general notion of band-limitedness; the considered bands are any subset Ω of the frequency domain W, not necessarily corresponding to frequencies inside an interval. In Section 3, Proposition 3.5.2 states that, in the graph setting, Ω-bandlimited signals can be reconstructed by sampling on $S \subset V$ whenever the N column vectors $\{\mathbb{1}_{V-S}|U_\Omega\}$ are linearly independent[1], that is, a *basis* of $\mathcal{V}$. In fact any set of column vectors U_Ω can be completed to a basis by adding suitable canonical basis vectors. In [4, 22] a different point of view on the same phenomenon is provided. A related randomized algorithm is described in [104], and, despite apparent loose pertinence, was motivational in the early stages of this project.

Natural choices for the smoothness models build upon the graph's adjacency. The formal analogue of a classical time-frequency ω-signal is a k-bandlimited graph signal whose k first Fourier coefficients are non-null [22, 4]. Irregular sampling of k-bandlimited graph signals has been studied first by Pesenson [80, 81] who introduced the notion of uniqueness set associated to the subspace of k-bandlimited graph signals. If two k-bandlimited graph signals are equal on a uniqueness set, they are necessarily equal on the whole graph. Random sampling was described in [86]; a probability measure on the vertexes can be built, using the graph laplacian, in such a way that sampling k-bandlimited signals, on a random subset of the vertexes, admits stable reconstruction with high probability. Several other sampling schemes exist in the literature, such as schemes

[1] or equivalently for $\{\mathbb{1}_S|U_{W-\Omega}\}$.

based on a bipartite decomposition of the graph [72], on a decomposition via maximum spanning trees [75, 95, 53].

On the other hand, it is useful to look at one example in which a graph signal sampling theorem can be formulated in the same way as the classical one. To this end, I recall the classical statement with a proof that is particularly insightful; see for example [51, page 177]. This gives the opportunity to formally introduce the concept of *Dirac comb*:

DEFINITION 2.2.1. *Given the Fourier transform $\mathscr{F}$, a tempered distribution* Ш *will be called Dirac comb if both* Ш *and* $\mathscr{F}\{\text{Ш}\}$ *have uniformly discrete support.*

The proof of the Shannon-Nyquist theorem that I present is based on the Poisson Summation Formula:

THEOREM 2.2.2 (Poisson Summation Formula). *As tempered distributions on* $\mathbb{R}$*:*

$$\mathscr{F}\Big\{\sum_{n\in\mathcal{Z}}\delta(x-nT)\Big\} = \frac{1}{T}\sum_{k\in\mathcal{Z}}\delta(\nu-k/T). \tag{85}$$

One immediately notices that $\text{Ш}_T = \sum_{n\in\mathcal{Z}}\delta(x-nT)$ is an example of Dirac comb. Furthermore $\text{supp}(\text{Ш}_T)$ is not only uniformly discrete, but also a subgroup of $(\mathbb{R};+)$.

THEOREM 2.2.3 (Shannon-Nyquist). *Let $f:\mathbb{R}\to\mathbb{R}$ be a Schwartz function*[2]*, such that $\mathscr{F}\{f\}$ is supported in $[-\omega_s/2,\omega_s/2]$. Then f can be exactly reconstructed from its values on $\frac{1}{\omega_s}\mathcal{Z}$.*

PROOF. Let the sampled version of f be represented by the distribution $f(x)\sum_{n\in\mathcal{Z}}\delta(x-n/\omega_s)$. Using Theorem 2.2.2:

$$\begin{aligned}\mathscr{F}\Big\{f(x)\sum_{n\in\mathcal{Z}}\delta(x-n/\omega_s)\Big\}(\nu) &= \mathscr{F}\{f\}(\nu) * \sum_{k\in\mathcal{Z}}\omega_s\delta(\nu-k\omega_s)\\ &= \omega_s\sum_{k\in\mathcal{Z}}\mathscr{F}\{f\}(\nu-k\omega_s).\end{aligned} \tag{86}$$

where $*$ denotes the convolution product. Therefore, because f is band-limited:

$$\mathbb{1}_{[-\omega_s/2,\omega_s/2]}(\nu)\mathscr{F}\Big\{f(x)\sum_{n\in\mathcal{Z}}\delta(x-n/\omega_s)\Big\}(\nu) = \mathscr{F}\{f\}(\nu).$$

□

Roughly speaking, the result is ensured by the existence of a sufficiently rich collection of Dirac combs, corresponding to the sampling frequency, and being an orbit with respect to the group of translations and the group of modulations. In fact, in some cases, similar collections can be constructed on graphs too. This is the case, for instance, when the graph at hand is a Cayley graph of a finite abelian group G. In this case, G has an action $g\in G, T_g:\mathcal{V}\to\mathcal{V}$ by translation, and the group of its characters $\hat{G}$ has an action $\chi\in\hat{G}, M_\chi:\mathcal{V}\to\mathcal{V}$ by modulation, as in the continuous case. Furthermore, whenever the degree $|G|$ is not prime, some functions behave like Dirac combs. These are functions whose support is a set of scattered isolated points, on both domains V and W. They will be essentially determined by the factors of $|G|$, that is,

[2]For the purpose of this introduction a stronger formulation is not necessary.

corresponding to each subgroup $H < G$. One such Dirac comb is simply the indicator function of each subgroup H (rescaled to have unit ℓ^2 norm):

$$\text{Ш}_H = \frac{1}{\sqrt{|H|}}\mathbb{1}_H(g), \quad U^*\text{Ш}_H = \frac{1}{\sqrt{|H^\perp|}}\mathbb{1}_{H^\perp}(g),$$

where $H^\perp$ denotes the subgroup of $\hat{G}$ of characters that are constant on H (see, for example [26]). The rest of the Dirac comb collection is given by the $T_G \times M_{\hat{G}}$-orbit of Ш_H. These are graph signals that have particularly small support on both the graph and frequency domains; in fact in this case all the graph signals $f \in \mathcal{V}$ reaching the smallest possible value of $\|f\|_0 + \|U^*f\|_0$ can be obtained in this way; this result is obtained by rephrasing the finite Fourier analysis results from [69, 100] in the graph signal case. This case is further elaborated in Section 5.

Let $J_V \subset V$ and $J_W \subset W$ be the supports of a graph signal $f \in \mathcal{V}$ and of its graph Fourier transform $\hat{f} := U^*f$ respectively. As a result the following linear combination of column vectors is null:

$$\sum_{i=1}^{N} f_i e_i - \sum_{i=1}^{N} \hat{f}_i u_i \equiv 0.$$

Therefore the set of column vectors $\{\mathbb{1}_{J_V}, U_{J_W}\}$ is linearly dependent.

What I have just observed is that the relevant information about sampling of graph signals is encoded in the linear independence/dependence relationships of some subsets of columns of the $N \times 2N$ matrix $[\mathbb{1}|U]$.

REMARK 2.2.4. *Sampling properties of band-limited signals only depend on the eignefunctions of $\mathcal{L}$, not on the spectrum. On the other hand, smoothness properties of these band-limited signals substantially rely on the spectrum.*

2.3. Uncertainty.

Efforts have recently been spent in literature towards understanding how the classical time-frequency *uncertainty principles* extend to the graph-frequency setting [79]; this is of course a wide topic, as localization in the graph and frequency domains can be defined in a variety of different ways. The problem is not considered well understood, albeit most studies agree in identifying, as main issue, the possibility for some u_k to be highly localized in the graph domain, depending on the topology of the graph [86].

The expression *uncertainty principle* seems at first to have a sinister connotation, as a principle showing that some things are impossible to make. On the contrary, the uncertainty principle is essentially the reason why stable recovery from a sample of signals or images can in fact be performed, despite a significant amount of missing information; see for example [33]. In particular, uncertainty principles can provide guarantees that if a signal has a sparse decomposition in a dictionary of incoherent atoms, this is indeed a unique representation that can be recovered via optimization [36]. This idea underlies the recent wave of sparse signal processing techniques, with applications such as denoising, source separation, inpainting, and compressive sensing.

Uncertainty principles such as the ones presented in [33, 36, 45, 21] are important examples. It is desirable that the dictionary atoms are jointly localized in time and frequency, and uncertainty principles characterize the resolution tradeoff between these two domains.

A first family of uncertainty principles involves a notion of measure of the support of a signal. Such notions can be either support measures counting the number of non-zero elements, or concentration measures, such as ℓ^p-norms. An important distinction is that these sparsity and concentration measures are not localization measures, as they disregard whether the support is clustered in a connected region of the vertex domain or not. A graph signal is nothing but a finite-dimensional vector; one example based on this idea is provided by the Elad-Bruckstein uncertainty principle [36]:

$$\frac{\|f\|_0 + \|U^* f\|_0}{2} \geq \sqrt{\|f\|_0 \|\hat{f}\|_0} \geq \frac{1}{\mu_G}$$

with $\mu_G = \max_{ik} |\langle \delta_i, u_k \rangle|$, *coherence*. Given $v \in \mathcal{V}$, having Fourier transform $U^* v$, the functional $\|v\|_0 = |\{i | v_i \neq 0\}|$ is called 0-norm in the literature (formally this is not a norm). An example of $\|\|_p$ based uncertainty principle for graph signals is [106]. These results typically depend on the mutual coherence between the graph Laplacian eigenvectors and the canonical basis of deltas. The bounds are global, so they might be dramatically affected by a bottleneck that is only caused by the edges of a few vertexes.

A second family of uncertainty principles involves a notion of distance. It goes along the line of the Heisenberg uncertainty principle in physics, see for example [39, 64]. In this case *uncertainty* corresponds to how much the signal spreads around some reference point, for example measuring the variance of the signal (seen as a measure) in each domain. The Heisenberg uncertainty principle states that the product of variances in the time and in the Fourier domains cannot be arbitrarily small. The generalization of this uncertainty principle to the graph setting is complex since there does not exist a simple formula for the mean value or the variance of graph signals, in either the vertex or the graph spectral domains. In [1] the geodesic graph distance $d(\cdot, \cdot)$ is used to constrain the spread of a graph signal around a given reference vertex:

$$\frac{f(i)}{\|f\|_2} \leq C e^{-k d(i,j)}, \text{ for some } j \in V$$

while on the frequency domain one might use: $\langle f, \mathcal{L} f \rangle / \|f\|^2$ which is a measure of smoothness. While there is in fact a tradeoff, the classical idea that a signal cannot be simultaneously localized in the time and frequency domains around *any* reference point does not always carry over to the graph setting.

REMARK 2.3.1. *Uncertainty principles based on a measure only depend on the eignefunctions of $\mathcal{L}$, not on the spectrum. On the other hand, uncertainty principles involving a notion of distance substantially rely on the spectrum.*

Local uncertainty principles were also studied, see [79]; these aim at a robust formulation that isn't drastically affected by a few bottlenecks in either of the domains.

2.4. Frames.

Once a definition of localization is chosen, a problem of great relevance in practice is finding frames of functions that are all jointly localized in the graph and frequency domains, possibly exhibiting a hierarchical organisation resembling that of wavelet frames; this has in fact been the main focus within the community, see for example [25, 57, 72, 48, 87, 35, 61, 75, 98, 97, 43]. The goal of these techniques is to sparsely represent different classes of graph signals or efficiently reveal important structural properties. Again, there is a little to no general theoretic approach

and many choices are empirically justified. In [48, 43] a form of *Littlewood-Paley decomposition* of the spectrum of $\mathcal{L}$ is attempted. Two main questions have to be addressed. The first is if and in which sense the resulting frame exhibits joint localization, while the second is how much the redundancy of the frame can be eliminated without loosing the main properties. These two questions have been addressed in [27] in the general setting of Dirichlet Spaces satisfying certain hypotheses, possibly having a rather small overlap with the scope of my work, as the hypotheses, translated in the graph signal processing setting, turn out to be rather artificial.

Motivated by this scenario, I'm currently investigating the functions exhibiting a certain type of joint localization, that I'm now going to describe.

Dirac combs in time-frequency analysis, as well as signals on Cayley graphs of some groups, play a pivotal role in the design of particularly well behaved set of signals, such as multi-scale frames and bankfilters in general. Indeed, in presence of a cascade of sampling theorems at different scales, corresponding to a cascade of nested bands in the frequency domain, one has the following sequence:

$$L^2(\mathbb{R}) \supset \cdots \supset V(B_n) \supset V(B_{n-1}) \supset \cdots \tag{87}$$

where, in the time-frequency case, one may have for example[3]:

$$B_n = [-\omega_n/2, \omega_n/2], \quad \omega_n = 2^n\omega, \quad \forall n \in \mathcal{Z}$$

and

$$V(B_n) = \text{span}\left\{T_{j/\omega_n}[\phi_n]\right\}_{j\in\mathcal{Z}}, \quad \text{supp}\left(\mathscr{F}[\phi_n]\right) \subseteq B_n,$$

for some bandlimited $\{\phi_n \in L^2(\mathbb{R})\}_{n\in\mathcal{Z}}$, and $T_x[f](\cdot) = f(\cdot - x)$.

A striking result of multi-resolution analysis is that one can choose ϕ_n in such a way that[4]:

- $\phi_n(\cdot) = D_{j/\omega_n}[\phi_1(\cdot)] \propto \phi_1(\cdot\ \frac{\omega_n}{j})$,
- there exists $\psi_1 \in L^2(\mathbb{R})$ allowing:

$$V(B_n) = W(B_n) \oplus V(B_{n-1}), \quad W(B_n) = \text{span}\left\{T_{j/\omega_n}[\psi_n]\right\}_{j\in\mathcal{Z}},$$

  with $\psi_n(\cdot) = D_{j/\omega_n}[\psi_1(\cdot)]$.

were the operator D, called dilation, does nothing but stretching the support, while maintaining the L^2-norm unaltered. There follows an orthogonal decomposition $L^2(\mathbb{R}) = \bigoplus_{n\in\mathcal{Z}} W(B_n)$ such that each $W(B_n)$ is band-limited, and obtained by transforming one function ψ_1 through T, D; this result had tremendous impact in the applications, see for example [65, 28]. An example is provided by $\phi_1 = \mathscr{F}^{-1}[\mathbb{1}_{B_1}], \psi_1 = \phi_1 - \phi_2$, called *Haar wavelets*. Formal similarity allows to apply the same machinery to the case of signal processing on Cayley graphs of abelian groups.

The following remark pertains the problem of more general graphs:

REMARK 2.4.1. *Roughly speaking, from the proof of Theorem 2.2.3 one can notice that, in order to establish the decomposition (87) everything one needs is: a sequence of nested bands* $\cdots \supset B_n \supset B_{n-1} \supset \ldots$ *and a rich enough collection of Dirac combs, ordered in nested families* $\cdots \supset \text{Ш}_n \supset \text{Ш}_{n-1} \supset \ldots$ *such that every Dirac comb in* Ш_n *intersects* B_n *in only one of its spikes, and every point in* B_n *is inside the support of one Dirac comb from* Ш_n. One should notice that, so far, Remark 2.3.1 still holds. *The eigenvalues only come into play as soon as each band* B_n *is*

[3] A more general formulation is not necessary for the purpose of this introduction.

[4] I omitted the normalization factor in the definition of dilation, in order to lighten the notation.

required to be neighbourhood of a reference frequency, or if spatial localization is imposed for the functions in $V(B_n)$, as this would rely on the adjacency notion.

An orthogonal decomposition is always possible, but its interpretation through translations and dilations crucially relies on the specific form that the Dirac combs take; in any non-standard scenario this is likely to get lost.

For functions on $\mathbb{R}^n$, non-standard Dirac combs, properly called *Fourier quasi-crystals*, have been studied extensively; see for example [58]. The corresponding possible generalizations of multi-resolution analysis have been explored as well [29]. Concerning graph signal processing, signals with small $\|\|_0$ support on both domains might be considered a generalization of Dirac combs. The situation is non-standard in a different sense as compared to $\mathbb{R}^n$, as rather than looking for discrete supports that don't have a group structure, the spatial domain itself has no group structure. The vast majority of the graph signal processing community is focused on building efficient algorithms for some specific purpose, and, with a few exceptions, comes from a background in analysis or engineering. I'm not aware of in depths studies in this sense, from a combinatorial point of view.

2.5. Combinatorial Framework.

To better understand the phenomena described above, I propose to consider the N-by-$2N$ matrix $[\mathbb{1}, U]$ as a point on the Grassmannian $\mathrm{Gr}_N^{2N} := \mathrm{Gr}\Big(N, \ell^2(V) \oplus \ell^2(W)\Big)$. The latter, as algebraic variety, is classically stratified in toric subvarieties called *Strata*, whose points exhibit the same combinatorial properties; each *Stratum* corresponds to a *Matroid* [78]. Given the set of indexes $[2N]$, a matroid is defined, in several ways, by a subset of the power-set $2^{[2N]}$. For example, a matroid $\mathcal{M}$, corresponding to a matrix $M \in \mathrm{Gr}_N^{2N}$, can be regarded as the collection of its *circuits* $\mathcal{C}_\mathcal{M} \subset 2^{[2N]}$, that is, minimal dependent sets. This means that for $C \in \mathcal{C}_\mathcal{M}$ the columns of M labelled by C are linearly dependent[5] and C is minimal with respect to this property. Another way is by $\mathcal{B}_\mathcal{M}$, the set of *bases*, maximal independent sets. Using this terminology, I can show that joint localization of graph signals, as in Proposition 3.5.3, boils down to the collection of circuits of the matroid $[\mathbb{1}, U]$. In a similar way, the problem of reconstructing band-limited signals from subsampling [86] is captured by the set of bases of the same matroid, as in Proposition 3.5.2. To the best of my knowledge this analogy has not been exploited so far in the literature, besides for the sporadic use of the term *girth* in Compressed Sensing. Furthermore, the construction of frames of graph signals with localization properties can benefit from Propositions 3.5.4 and 3.5.5.

The matrices of the form $[\mathbb{1}, U]$, with U being unitary, do not cover the entire $\mathrm{Gr}_N^{2N}(\mathbb{C})$, but describe a sub-variety. The latter inherits a stratification from $\mathrm{Gr}_N^{2N}(\mathbb{C})$. I can show that the corresponding matroids admit a combinatorial characterization as *identically self-dual* matroids (ISD). These are matroids with several properties, especially regarding the lattice theoretic concepts of *flats* and *cycles*, the former being intersections of hyperplanes, the latter being unions of circuits. It turns out that, in a ISD matroid, the lattice of flats is dual to the lattice of cycles; this makes particularly relevant the so called *cyclic flats*, that by definition are both flats and cycles. Structural properties of ISD matroids are reviewed in Section 4, also offering some possibly novel results.

[5] Throughout, linear dependence is meant on $\mathbb{C}$.

The main technically contribution is an in depth description of the special case of Cayley graphs of finite abelian groups, in which the Dirac combs play a pivotal role; see Section 5. This case corresponds to a class of representable ISD matroids, which I call *Abelian Matroids*. They are highly symmetrical objects, exhibiting a striking amount of properties. In particular, under one hypothesis, the lattice of cyclic flats is atomistic, and in fact is the intersection between the union lattice generated by the atoms, and the intersection lattice of their complements (Theorems 5.2.7 and 5.2.11 and Corollary 5.2.12).

Based on this insight, my present effort is to understand how such scenario translates in the case of a general ISD matroid, obtained form a generic graph Fourier transform U, using structural results about matroids their cyclic flats [**15**, **34**].

So far I only focused on the operator U regardless the graph structure. Within this project, another line of research pertains the spectrum of the operator $\mathcal{L}$; this is covered in Section 6, mostly conjectural at the moment. It includes a result about the existence of a nontrivial spectrum such that U corresponds to an admissible $\mathcal{L}$ (Proposition 6.1.2) and an open problem about a second stratification of the group $U(N)$ through admissible spectra (Remark 6.2.3), possibly related to the matroidal stratification described above.

3. Sparsity and Matroids

In this section I present the mathematical problem corresponding to the framework outlined in the previous section. A few known results are reported here, in order to provide a self-consistent foundation. The main results of this section are Theorem 3.2.2, Proposition 3.5.4 and Proposition 3.5.5, and, despite simple, seem to be novel in this context. Roughly speaking they describe how, in principle, a combinatorial object, an identically self-dual matroid constructed from the eigenfunctions of the graph laplacian operator, gives hints towards defining a frame or basis of highly localized (Dirac-comb-like) graph signals. A prototypical example of such collection is the collection of Dirac Combs, as signals on a Cayley graph of a finite abelian group. An in depth combinatorial description of the latter is postponed to Section 5.

3.1. Graph Signals.

A graph will be referred to as $\{V, A\}$, where $V := [N]$ is the ordered *vertex set* (the order is arbitrary but fixed), while $A \in M_{V\times V}(\mathbb{R})$ is the *edge set*, being A_{ij} weight of the edge between vertexes $i, j \in V$; zero weights are considered missing edges. The graph is considered *undirected*, that is $A^T = A$, and *acyclic* that is $A_{ii} = 0, \forall i \in V$. Every vertex has *degree* $d_i = \sum_{j\in V} A_{ij}$.

I call *signals* the elements of $\mathcal{V} := \ell^2(V) \equiv \mathbb{C}^V$, with $\langle\cdot,\cdot\rangle$ canonical hermitian inner product, and $\|\cdot\|$ the ℓ^2 norm. I refer to the *canonical basis* of $\mathcal{V}$ as $\{e_i\}_{i\in V}$. Given a signal $v \in \mathcal{V}$, its components are $v(i) := \langle v, e_i\rangle$. A matrix $O \in M_{V\times V}(\mathbb{C})$ is a linear operator on $\mathcal{V}$, with matrix entries O_{ij}, and O^* its adjoint.

The operator $\mathcal{L} = D - A$ is called Graph Laplacian, where $D = \operatorname{diag}\big(d_1, ..., d_N\big)$. The definition implies $\mathcal{L}$ is a real symmetric matrix, diagonally dominant, with nonpositive off-diagonal entries and positive diagonal entries. As a result it has a spectrum of real nonnegative eigenvalues, corresponding to orthogonal eigenspaces; see for example [18].

I choose an order on the spectrum (eigenvalues with multiplicity are repeated accordingly) $\{\lambda_i\}_{i\in W}$, with $W = [N]$, such that $\lambda_i \leq \lambda_j$ whenever $i < j$, and an orthonormal basis of eigenvectors $\{u_i\}_{i\in W}$. W will be called *frequency domain*. I introduce the space of vectors $\mathcal{W} := \ell^2(W)$, with the canonical basis $\{f_i\}_{i\in W}$.

Remark 3.1.1. *In the case an eigenvalue has multiplicity higher than one, that is $\lambda_j = \cdots = \lambda_{j+k}$, the choice of $\{u_i\}_{j\leq i\leq j+k}$ is not unique. This problem is shared by other algebraic approaches to graph signal processing, see for example* [101]. *This thesis does not address the problem, as I mainly focus on the properties of sets of orthonormal signals $\{u_i\}_{i\in W}$ rather than of the corresponding laplacian operator. Relevance of the issue remains unclear, in particular concerning the discussion in Section 6. A reasonable ansatz is choosing $\{u_i\}_{j\leq i\leq j+k}$ with minimal supports*[6].

For all signals $v \in \mathcal{V}$ I call *graph Fourier transform* the linear operator $U^* : \mathcal{V} \to \mathcal{W}$, with U a unitary matrix having $\{u_i\}_{i\in W}$ as columns, so that $u_i = U f_i$. The operator $U : \mathcal{W} \to \mathcal{V}$ is called anti-Fourier transform.

I introduce the auxiliary space $\mathcal{X} := \mathcal{V} \oplus \mathcal{W}$. I use the notation $E := V \cup W = \{1, ..., N, N + 1, ..., 2N\}$. Unless otherwise stated, it is meant that $i \leq N$ points to $i \in V$, while $i > N$ points to

[6]The matroids corresponding to different choices are related anyways, one being weak image of the other.

$i - N = j \in W$. The canonical basis of $\mathcal{X}$ is referred to as $\{g_i\}_{i\in V\cup W}$ that is $\{e_i \oplus 0\}_{i\in V} \cup \{0 \oplus f_i\}_{i\in W}$.

Signals on $\mathcal{V}$ are considered embedded in $\mathcal{X}$ via $J : \mathcal{V} \to \mathcal{X}$ such that $v \mapsto v \oplus U^*v$. The image $J(\mathcal{V})$ is spanned by the orthonormal basis $\{\frac{\sqrt{2}}{2}(e_i \oplus U^*e_i)\}_{i\in V}$. This linear subspace will be denoted by $\mathcal{S}_U \subset \mathcal{X}$.

I denote with $P_{\mathcal{S}} : \mathcal{X} \to \mathcal{X}$ the usual orthogonal projection on a linear subspace $\mathcal{S} \subseteq \mathcal{X}$, while the application $\Pi_{\mathcal{S}} : \mathcal{X} \to \mathcal{X}/\mathcal{S}^{\perp}$ is the canonical projection to the quotient space. In the cases in which $\mathcal{S}$ is a coordinate space, the representatives of the quotient are chosen by disregarding the coordinates that are identically zero in $\mathcal{S}$.

The function $\|\|_0$ is defined in $\mathcal{X}$ as $\|x\|_0 := \#\{i \,|\, x(i) \neq 0\}$. Analogous definitions hold in $\mathcal{V}$ and $\mathcal{W}$, so that $\|v \oplus U^*v\|_0 = \|v\|_0 + \|U^*v\|_0$.

3.2. Sparsity Classes.

I shall study the following distinguished subspaces of $\mathcal{V}$:

Definition 3.2.1.

$$\mathcal{D}_s := \{v \in \mathcal{V} \,|\, \|v \oplus U^*v\|_0 = s\}$$

Clearly $\mathcal{D}_{2N} = \mathcal{V}$ and $\mathcal{D}_0 = \{0\}$. Besides these trivial cases, $\mathcal{D}_s$ is not in general a linear subspace of $\mathcal{V}$, but rather a union of linear subspaces. Referring to multi-indexes $\{i_1, ..., i_s\} =: J \subseteq V \cup W$, I have:

$$\mathcal{D}_s = \bigcup_{|J|=s} \mathcal{D}_J$$

given by vectors $v \in \mathcal{V}$ such that $v \oplus U^*v$ is supported on J. The classes form a Poset with respect to the partial order given by the inclusion $\subseteq$. Excluding $\{0\}$ it may have several minimal elements, that are classes not including any other non-empty classes, while there is only one maximal element, being $\mathcal{D}_{2N}$. I shall discuss the importance of such minimal non-trivial elements later.

The following few results introduce the combinatorics of representable matroids, as applied to the classes $\mathcal{D}_J$.

Theorem 3.2.2. *For any multi-index J, let P_J be the orthogonal projection on $\mathcal{S}_J := \text{span}\{g_i\}_{i\in J}$. Consider the following $2N \times 2N$ matrix:*

$$O = \begin{bmatrix} -\mathbb{1} & U \\ U^* & -\mathbb{1} \end{bmatrix},$$

where $\mathbb{1}$ is the $N \times N$ identity matrix. Then I have:

$$\mathcal{D}_J = \Pi_{\mathcal{V}} P_J \ker\left(P_J O P_J\right) \tag{88}$$

Proof. Let P_U denote the orthogonal projection on $\mathcal{S}_U$ and P_I on $\mathcal{S}_J$. Clearly $\mathcal{D}_J = \Pi_{\mathcal{V}}(S_J \cap S_U)$. It is shown in [10] that the eigenspace of eigenvalue 1 of the product $P_J P_U$ is exaclty $S_J \cap S_U$ (and by symmetry the same holds for $P_U P_J$). It follows:

$$S_J \cap S_U = P_J \ker(P_J P_U - \mathbb{1}) = P_J \ker(P_U P_J - \mathbb{1}) = P_J \ker(P_J P_U P_J - \mathbb{1})$$
$$= P_J \ker(P_J(P_U - \mathbb{1})P_J).$$

It remains to provide en expression for P_U.

The operator $\begin{bmatrix} 0 & U \\ U^* & 0 \end{bmatrix}$ is unitary, self-adjoint and traceless at the same time, and as a result has eigenvalues $+1, -1$ both with multiplicity N; the corresponding eigenspaces are $\mathcal{S}_U = \text{span}\{\frac{\sqrt{2}}{2}(e_i \oplus U^* e_i)\}_{i\in V}$ and $\mathcal{S}_U^\perp = \text{span}\{\frac{\sqrt{2}}{2}(e_i \oplus -U^* e_i)\}_{i\in V}$ respectively. By adding the identity operator, I get an operator with eigenvalues $+2, 0$, and same eigenspaces. It follows that:

$$P_U = \frac{1}{2}\Big(1 + \begin{bmatrix} 0 & U \\ U^* & 0 \end{bmatrix}\Big)$$

so that $O = -2P_U^\perp = 2(P_U - 1)$, and the claim follows.

□

Let A_{IJ} denote the $|I| \times |J|$-matrix obtained from A by erasing columns not labelled by J and rows not labelled by I.

COROLLARY 3.2.3. *The class $\mathcal{D}_J$ is non empty if and only if the following holds:*

$$\det(O_{JJ}) = 0 \tag{89}$$

or, equivalently, the columns of O labelled by J are linearly dependent. In fact, this is also equivalent to the linear dependence of the columns labelled by J of the matrix:

$$[\mathbb{1}|U] \tag{90}$$

The last statement justifies the use of the letter $\mathcal{D}$, as *Dependent Sets*. Roughly speaking, linear dependence of the columns labelled by J means that, there is a null linear combination of the vectors $\{e_i\}_{i\in J\cap V}, \{u_j\}_{j\in J\cap W}$ with non-zero coefficients; in that case the coefficients of the $\{e_i\}_{i\in J\cap V}$ are nothing but the values of a function in $\mathcal{D}_J \subset \mathcal{V}$, while the coefficients of the $\{u_j\}_{j\in J\cap W}$ provide the values of its Fourier transform.

Just like in classical time-frequency analysis, sparsity, as a measure of localization on the two domains, is bounded from below. As mentioned in Section 2 results formalizing this constraint are called *uncertainty principles*. One such result is presented in Section 3.5.

I already mentioned the ordering of the sparsity classes $\mathcal{D}_J$ by inclusion; from a combinatorial point of view, $\mathcal{D}_J$ is particularly relevant if it is non-empty and minimal:

DEFINITION 3.2.4. *A multi-index $J \subseteq [2N]$ is called a circuit if D_J is non empty and there is no support $J' \subset J$ admitting a non-empty $\mathcal{D}_{J'}$. The collection of circuits is denoted by $\mathcal{C}_U$.*

The combinatorial object described by $\mathcal{C}_U$ admits a number of equivalent descriptions, all of interest from a different point of view.

3.3. Rank functions and Bases.

Here an alternative description is offered, introducing the notion of combinatorial rank, from the algebraic counterpart. I follow here the historical viewpoint of [40]. A more concise language is adopted later.

The space $\mathcal{S}_U$ (as well as its orthogonal complement $\mathcal{S}_U^\perp$) can be described by a *configuration*, that is a $C_U \in M_{N\times 2N}(\mathbb{C})$ rank-N-matrix (resp. $C_U^\perp$). One can interpret this matrix as either a

collection of N vectors in $\mathcal{X}$ that span $\mathcal{S}_U$ (resp. $\mathcal{S}_U^\perp$), or as a collection of $2N$ vectors in $\mathcal{X}/\mathcal{S}_U^\perp$ (resp. $\mathcal{X}/\mathcal{S}_U$) that redundantly span such space. Namely:

$$\begin{aligned} &C_U : \mathcal{X} \to \mathcal{X}/\mathcal{S}_U^\perp, \quad g_i \mapsto C_U g_i, i \in [2N], \quad \mathcal{S}_U^\perp = \ker(C_U) \\ &C_U^\perp : \mathcal{X} \to \mathcal{X}/\mathcal{S}_U, \quad g_i \mapsto C_U^\perp g_i, i \in [2N], \quad \mathcal{S}_U = \ker(C_U^\perp). \end{aligned} \tag{91}$$

Remark 3.3.1. *The choice of configuration is clearly not unique, as the action of $GL(N)$ by left multiplication does not modify the subspace represented by a configuration. Instead, the action of $GL(2N)$ by right multiplication does modify the subspace represented by a configuration;*

Definition 3.3.2. *Given any $J \subset \{1, ..., 2N\}$, the rank functions $r, r^\perp, s$ are defined by:*

$$\begin{aligned} r(J) &:= \dim(\text{span}\{C_U g_i\}_{i\in J}), \\ r^\perp(J) &:= \dim(\text{span}\{C_U^\perp g_i\}_{i\in J}), \\ s(J) &:= \dim(\mathcal{D}_J). \end{aligned} \tag{92}$$

Equivalent definitions can be found in literature: $r^\perp(J) = \dim(\mathcal{S}_U^\perp/(\mathcal{S}_U^\perp \cap \mathbb{C}^{\bar{J}})) = \dim(\mathbb{C}^J/(\mathbb{C}^J \cap \mathcal{S}_U))$, where $\bar{J} = [2N] - J$, and equivalently $r(J) = \dim(\mathcal{S}_U/(\mathcal{S}_U \cap \mathbb{C}^{\bar{J}})) = \dim(\mathbb{C}^J/(\mathbb{C}^J \cap \mathcal{S}_U^\perp))$.

Lemma 3.3.3. *For any $J \subseteq [2N]$, with $\bar{J} = [2N] - J$, the following holds:*

$$r^\perp(J) = |J| - s(J), \quad r(J) = N - s(\bar{J}) \tag{93}$$

Corollary 3.3.4. *The class $\mathcal{D}_J$ is non-empty if and only if $r^\perp(J) < |J|$.*

The rank functions $r^\perp, r$ (or, in their modern combinatorial notation $\text{rank}, \text{rank}^*$) play a pivotal role in the theory of matroids. So far, the information that U is unitary has not been used. This in fact has a remarkable consequence for the rank functions:

Proposition 3.3.5. *In the context of this section, for any $J \subseteq [2N]$ one has $r^\perp(J) = r(J)$.*

From now on, only the notation $\text{rank}(J)$ will be used for the rank function.

I described how the combinatorial information contained in the matrix $[\mathbb{1}|U]$ is encoded in the function $\text{rank}(\cdot)$, and how this in turn determines the sparsity classes $\mathcal{D}_J$, and the corresponding spaces of functions through Theorem 3.2.2.

The combinatorial information stored in the function $\text{rank}(\cdot)$ can be alternatively encoded in a number of equivalent ways. An important one is the following:

Definition 3.3.6. *The multi-indexes of the following collection:*

$$\mathcal{B}_U := \{J \subset [2N] \,|\, |J| = \text{rank}(J) = N\}$$

are called bases.

The terminology is self-explanatory, the columns of $[\mathbb{1}|U]$ corresponding to a multi-index $J \in \mathcal{B}_U$ are linearly independent. The function rank is completely determined by the corresponding collection of bases $\mathcal{B}_U$. In particular :

Lemma 3.3.7.

$$\text{rank}(J) = \max_{I\in\mathcal{B}_U} |I \cap J|$$

From now on, whenever I talk about a *Matroid*, I refer, equivalently, to its rank function or to its collection of bases. The collections of circuits $\mathcal{C}_U$ and of bases $\mathcal{B}_U$ can be obtained one from the other, see for example [78].

3.4. Grassmannian Strata and Polytopes.

I already mentioned that the correspondence between matrix $[\mathbb{1}|U]$ and the corresponding matroid is not 1-to-1; this observation is simple but deep. Given a field $\mathbb{F}$, consider $\mathbb{F}^N$ and a linear subspace represented as row-space of a matrix M. The matrix can indeed be transformed in various ways and still correspond to the same matroid; see for example [16]:

PROPOSITION 3.4.1. *Let a matroid be represented by the $d \times N$ matrix M in $\mathbb{F}$,*

1. *the action of $GL(k, \mathbb{F})$ on M by left multiplication does not modify the row space, hence the matroid,*
2. *the action of $GL(N, \mathbb{F})$ on M by right multiplication does modify the row space, but, the Cartan subgroup $H < GL(N, \mathbb{F})$ given by the diagonal matrices, acting by right multiplication, does not modify the matroid.*
3. *the action of a field automorphism of $\mathbb{F}$ on M does not modify the matroid.*

Since I work on $\mathbb{C}$ the only relevant field automorphism is the complex conjugation.

Point *(2)* of Proposition 3.4.1 suggests to look at the orbit $H \cdot M$; its Zarisky closure $\overline{H \cdot M}$ is a toric sub-variety of the Grassmannian $\mathrm{G}_k(\mathbb{F}^N)$. This is in fact the algebraic object that corresponds to the matroid, see for example [40].

Going back to the setting of the previous section, in $\mathcal{X} \equiv \mathbb{C}^{2N}$ let $\mathrm{G}_N(\mathcal{X})$ denote the set of all N-dimensional linear subspaces of $\mathcal{X}$, the Grassmannian $\mathrm{G}_N(\mathcal{X})$. A classical way to describe Grassmannians as smooth manifolds is by defining the Plücker coordinates.

DEFINITION 3.4.2. *Given a N-dimensional linear subspace $\mathcal{S} := \mathrm{span}\{x_i\}_{i=1,\dots,d} \subset \mathcal{X}$, its Plücker coordinates are labelled by the multi-indexes $J \subset [2N]$ with $|J| = N$, and are given by:*

$$p_J(\mathcal{S}) := \det\left(\Pi_J x_1, \dots, \Pi_J x_N\right)$$

In the cases of my interest I consider the subspace $\mathcal{S}_U$: the vectors $\{x_i\}_{i=1,\dots,N}$ are any row-configuration, such as $\mathrm{span}\{\frac{\sqrt{2}}{2}(e_i \oplus U^* e_i)\}_{i \in V}$.

REMARK 3.4.3. *I have $p_J \neq 0$ if and only if $J \in \mathcal{B}_U$.*

These coordinates are defined up to a multiplicative constant, and they only depend on the subspace, not on the configuration.

DEFINITION 3.4.4. *For every multi-index J, consider $\delta_J \in \mathbb{R}^{2N}$ having components $+1$ for labels in J and 0 otherwise. The function $\mu : G_N(\mathcal{X}) \to \mathbb{R}^{2N}$ taking the following values:*

$$\mu(\mathcal{S}) := \frac{\sum_{J \in B_N} |p_J(\mathcal{S})|^2 \delta_J}{\sum_{J \in B_N} |p_J(\mathcal{S})|^2}$$

this is called Moment Map.

The image $\mu\big(\mathrm{G}_N(\mathcal{X})\big)$ will be the convex hull of all the vectors $\delta_J, |J| = N$. This is a convex polytope depending only on N and $2N$. Given a subspace $\mathcal{S}_U$ as above, there corresponds a collection of bases $\mathcal{B}_U$, which in turn corresponds to a set of vertexes $\delta_J, J \in \mathcal{B}_U$.

THEOREM 3.4.5 ([16]). *Consider a matroid represented by a matrix M. The image $\mu\big(\overline{H \cdot M}\big)$ is the convex hull of the vectors $\delta_J, J \in \mathcal{B}_M$.*

Roughly speaking, two matrices M_1, M_2 represent the same matroid if and only if they share the same moment map image of their Zarisky closed H-orbit.

3.5. Traces of Combinatorics of Graph Signal Processing.

In this section I introduce a few results that give an idea of how graph signal processing problems might benefit from combinatorics. The common tool provided by the matroid represented by $[\mathbb{1}|U]$ offers a common language for questions that are usually dealt with separately, such as uncertainty principles, sampling theory and bankfilters design. None of the results presented have direct practical application; they are mostly meant to provide an intuition. Some will be used, later in this thesis, to prove other results.

3.5.1. *Sampling.*

DEFINITION 3.5.1. *Given a subset $B \subset W$, which I call band, a function $v \in \mathcal{V}$ will be called B-bandlimited if* $\mathrm{supp}(U^* v) \subseteq B$.

The following result summarizes graph signal sampling theory, in its most simple form, as anticipated in Section 2, by mimicking the time-frequency results about non-lossy sampling. The result is essentially not new, but its formulation presented here takes advantage of the formalism introduced.

PROPOSITION 3.5.2. *Let $J \in \mathcal{B}_U$ be a multi-index from the collection of bases. Then any J_W-bandlimited function v is uniquely determined by its values $v(\bar{J}_V)$, where $J_W = J \cap W$ and $\bar{J}_V = V - (J \cap V)$.*

For any band $B \subseteq W$ there exists a multi-index $J \in \mathcal{B}_U$ such that $J \cap W = B$.

Clearly $|J_W| = |\bar{J}_V|$.

PROOF. Since the multi-index J is a basis, that is $\mathcal{S}_U \cap \mathbb{C}^{\bar{J}} = \{0\}$, it follows that the subspace $\mathcal{S}_U$ can be expressed as graph of a function $f : \mathbb{C}^J \to \mathbb{C}^{\bar{J}}$.

Again, because J is a basis, the columns of C_U labelled by J are linearly independent for any equivalent choice of configuration C_U. Let me choose the following configuration (written by columns), whose restriction to J is the identity:

$$\{\{e_i\}_{i\in J_V}, \{Uf_i\}_{i\in J_W}, \{U^* e_i\}_{i\in J_V}, \{f_i\}_{i\in J_W}\}$$

Since I always have $V \in \mathcal{B}_U$, the set $\{\{e_i\}_{i\in J_V}, \{Uf_i\}_{i\in J_W}\}$ will be a (non-orthogonal) basis of $\mathcal{V}$. But then I notice that $\{Uf_i\}_{i\in J_W}$ is a basis for the space of L_W-bandlimited functions. Since $\{e_i\}_{i\in \bar{J}_V}$ is the orthogonal complement of the remaining part of the basis (i.e. $\{e_i\}_{i\in J_V}$), it follows

that the orthogonal projection of L_W-bandlimited functions onto $\{e_i\}_{i\in\bar{J}_V}$ must be injective; by dimensional counting, this projection is also surjective, thus the first statement holds.

The final statement holds since the columns of U are orthonormal, so that any set of columns can be completed to a basis of $\mathcal{V}$ by adding suitable vectors from another basis (namely from $\{e_i\}_{i\in V}$). □

3.5.2. *Uncertainty Priciples.*

Referring to the sparsity classes $\mathcal{D}_J$, I anticipated that they are expected to be empty up to a certain threshold size, where a *circuit* J corresponds to a minimal non-empty sparsity class. The smallest size of a circuit of a matroid is called *girth*. The following result from [36] provides a lower bound from this value in a particularly simple form:

PROPOSITION 3.5.3.

$$\frac{\|v \oplus U^*v\|_0}{2} \geq \sqrt{\|v\|_0\|U^*v\|_0} \geq \frac{1}{\mu} \tag{94}$$

where $\mu := \max_{i,j} |\langle u_i, e_j\rangle|$.

PROOF. If a vector v has $\|v\|_2 = 1$ and is supported on K components, then $\|v\|_1 \leq \sqrt{K}$, and equality is attained when each of the K components has absolute value $1/\sqrt{K}$. So, given $v = \sum_i a_i e_i = \sum_j b_j u_j$, it follows:

$$1 = \langle v, v\rangle = \langle \sum_i a_i e_i, \sum_j b_j u_j\rangle = \sum_{i,j} a_i b_j \langle e_i, u_j\rangle$$

$$\leq \max_{i,j} |\langle u_i, e_j\rangle| \sum_i |a_i| \sum_j |b_j| \leq \mu\sqrt{\|v\|_0}\sqrt{\|U^*v\|_0}.$$

The result follows since $\frac{\|v\oplus U^*v\|_0}{2}$ and $\sqrt{\|v\|_0\|U^*v\|_0}$ are arithmetic and geometric means respectively.

□

This result, despite simple, is somewhat counterintuitive: a quantity of combinatorial nature, such as the girth of a matroid, is constrained by the maximum value of the matrix U from a representation, which roughly speaking is the cosine of an angle. This is a typical scenario in Compressed Sensing, see for example [6]. In one of the subsequent sections I will analyze in details a class of matroids for which this lower bound is attained. Besides providing a lower bound for the girth, [6] provides conditions under which a search for minimal supports can be performed by convex optimization on the ℓ^1 norm, which is highly more efficient to handle than a combinatorial search. Despite the affinity, the reader will notice that none of the combinatorial terminology is ever mentioned therein.

3.5.3. *Frames.*

Lastly, concerning the problem of bankfilters design, the following two results suggest how one can, in principle, build frames or bases of highly localized (Dirac-comb-like) functions on a graph, by using circuits of the matroid $[\mathbb{1}|U]$.

By definition, any circuit $C \subset [2N]$ has $|C| - \text{rank}(C) = 1$, that is $\dim(\mathcal{D}_C) = 1$. Therefore, modulo an overall complex phase, there is only one function supported on $C \cap V$ that has unit ℓ^2 norm, with graph Fourier transform supported on $C \cap W$. This function will be denoted as $f[C] \in \mathcal{V}$. Let $\hat{f}[C]$ denote the graph Fourier transform of $f[C]$. The function also corresponds to a vector in $\mathbb{C}^E$ whose coordinates on $\mathbb{C}^V$ are the values of $f[C]$ and whose coordinates on $\mathbb{C}^W$ are the values of $\hat{f}[C]$.

PROPOSITION 3.5.4. *Consider a representable ISD matroid represented by $[\mathbb{1}|U]$. Let $f[C]$ be the unit ℓ^2-norm element of $\mathcal{V}$ associated with a circuit C.*

For any distinct circuits C_1, C_2, the functions $f[C_1], f[C_2]$ are linearly independent.

Let $\{C_i\}_{i=1,\dots,k}$ circuits with $S := \bigcap_{i=1}^k C_i$, and $B \in \mathcal{B}(M)$ such that $S \cap B = C_j \cap B, \forall j$. Then $\{f[C_i]\}_{i=1,\dots,k}$ as vectors of $\mathbb{C}^E$ have linearly independent restrictions to vectors in $\mathbb{C}^{S \cap B}$.

PROOF. By definition both C_1, C_2 have proper intersections with both V and W, and since are distinct, either $C_1 \cap V \neq C_2 \cap V$ or $C_1 \cap W \neq C_2 \cap W$ or both. Assume the first holds. Then $f[C_1], f[C_2]$ are linearly independent because on $(C_1 - C_2) \cap V$ or $(C_2 - C_1) \cap V$ only one of them has non-zero values. The same argument holds in case $C_1 \cap W \neq C_2 \cap W$.

Since the matroid is ISD, given a basis $B \in \mathcal{B}$ and a function $f \in \mathcal{V}$ as above, the values of f on $E - B$ are obtained from its values on B by applying a bijective linear map $T_B : \mathbb{C}^B \to \mathbb{C}^{E-B}$. Consider the vectors in $\mathbb{C}^E$ corresponding to $\{f[C_i]\}_{i=1,\dots,k}$. Their restrictions to $\mathbb{C}^{E-B}$ have incomparable supports, so they are linearly independent. Because T_B is bijective, the image vectors in $\mathbb{C}^B$ are also linearly independent, and by construction they are supported on $S \cap B$.

□

PROPOSITION 3.5.5. *Consider an ISD matroid on $[2N] \equiv V \cup W$ represented by $[\mathbb{1}|U]$, with U unitary. Let $f[C]$ be the unit ℓ^2-norm element of $\mathcal{V}$ associated with a circuit C. For two distinct circuits C_1, C_2 let $S = C_1 \cap C_2$. If $S \subseteq V$ or $S \subseteq W$, then $f[C_1], f[C_2]$ are orthonormal as elements of $\mathcal{V}$; moreover, restricted as vectors of $\mathbb{C}^S$, they are orthogonal.*

PROOF. Assume $S \subseteq V$. The scalar product reads:

$$\langle f[C_1], f[C_2] \rangle_{\mathcal{V}} = \sum_{i \in S} f[C_1](i) f^*[C_2](i) = \sum_{j \in W} \hat{f}[C_1](j) \hat{f}^*[C_2](j) = 0,$$

where the first and the last equality hold because the common support of the two signals is S on V and $\varnothing$ on W. If $S \subseteq W$ a similar argument holds.

□

In fact, a set of N circuits determines an orthonormal basis of $\mathcal{V}$ whenever they only intersect, pair-wise, on V or on W.

4. Identically Self-Dual Matroids

Identically Self-Dual Matroids are a rich class of matroids exhibiting a strong symmetry, the symmetry by complementation; these matroids are required by the theory outlined in the previous sections. They come across as soon as Duality is introduced in Matroid Theory, and therefore have been known and studied for long. Despite that, I'm not aware of published comprehensive accounts entirely dedicated to their notable structure. In this chapter I review many properties of Identically Self-Dual Matroids, of interest for the purpose of this thesis. Despite most of the material is more or less well known, the overall path, as well as some observations, might in fact be novel.

I already introduced Linear Matroids from the stratification of Grassmannian varieties, and anticipated how their combinatorial axiomatization leads to general Matroids. Following the historical path, the axiomatization usually starts by defining Independent sets, in analogy to linear independence in vector spaces. Once the point of view is shifted to the combinatorial setting of general matroids, the axioms of independent sets are just one of a number of equivalent axiomatizations, each based on a different combinatorial object, used to describe matroids; Independent sets, Dependent sets, Bases, Circuits, Bonds, Flats, Hyperplanes, Cycles, Rank function, and Closure operator are the most common ones, each being more convenient in a different context. Each of these can be independently axiomatized, or can be obtained by either of the others [109]. Because this chapter is concerned with Duality, and I particularly care about small circuits for application purposes, I'm going to focus on the axiomatizations that are better behaved to both these respects at the same time: namely the one based on the Bases, and the one based on the so called Cyclic Flats.

By exchanging opinions within the Matroid Theory community, I came across the unwritten belief that the class of ISD matroids is structurally *just about as complex* as the class of general matroids. I first investigated this idea hoping to build tools useful for my problem. One of the results I present shows a way to injectively map ISD matroids, satisfying one condition, onto a simpler combinatorial structure. It is not clear yet how restrictive is the condition, but it seems quite clear that the image of this map is a rather small subset of the codomain.

Lastly I add a miscellanea of observations concerning the class of Representable ISD Matroids. Despite the fact that, for the application at hand, I'm particularly interested in representable rather than more general ISD Matroids, the problem of representability is notoriously unmanageable with combinatorial tools [68], and therefore beyond the scope of this thesis.

4.1. Preliminaries.

From the point of view of the matrix representation, bases are maximal subsets of columns that are linearly independent. Regardless, modern combinatorics define matroids, and their bases, as abstract combinatorial objects; the representable case turns out to be a special case. Even though I'm mostly interested in the axiomatization of matroids through cyclic flats, I present the axioms for bases first, to have a reference for practical convenience.

Definition 4.1.1. *A collection of subsets of $E = [N]$ is the collection of bases $\mathcal{B}(M)$ of a matroid M if:*

(B1) *$\mathcal{B}(M)$ is non-empty, and*

(B2) *for every* $B_1, B_2 \in \mathcal{B}(M)$*, and* $x \in B_1 - B_2$*, there is an* $y \in B_2 - B_1$ *such that* $(B_1 - x) \cup y \in \mathcal{B}(M)$.

It can be shown that axiom *(B2)* is a combinatorial version of the Laplace expansion of determinants [109, Chapter 4]. An immediate property of bases is that they all have the same size, equal to the rank of M. A subset of E is independent if and only if it is a subset of a basis.

I previously mentioned the concept of circuit of a representable matroid, as a minimal subset of the columns that are linearly dependent. The combinatorial counterpart follows:

DEFINITION 4.1.2. *The collection of circuits* $\mathcal{C}(M)$ *of a matroid* M *is the collection of minimal subsets of* E *such that they are not contained in any basis.*

An immediate consequence is that any proper subset of a circuit is an independent set. Furthermore, they are clearly all pair-wise incomparable sets.

A quite transparent relationship between circuits and bases comes from the following definition:

DEFINITION 4.1.3. *Given a basis* $B \in \mathcal{B}(M)$ *and* $e \in E - B$*, the fundamental circuit* $C(B, e)$ *is the unique circuit contained in* $B \cup e$.

In fact any circuit is fundamental with respect to some (B, e). Morover, the following property holds:

PROPOSITION 4.1.4. *Given a basis* $B \in \mathcal{B}(M)$*,* $f \in B$ *and* $e \in E - B$*, the set* $(B \cup e) - f$ *is a basis of* M *if and only if* $f \in C(B, e)$.

Roughly speaking, the fundamental circuits of a basis encode the exchange properties of that basis.

Duality is very easy to describe in terms of the bases:

PROPOSITION 4.1.5. *The collection* $\mathcal{B}(M^*) = \{E - B \mid B \in \mathcal{B}(M)\}$ *is the collection of bases of a matroid* M^*.

The matroid M^* is called *dual matroid* of M. This section is focused on matroids such that $M = M^*$, that is, *identically self-dual matroids*. Clearly this imply $E = [2N]$ and $\text{rank}(M) = N$.

The phenomenon described in Proposition 4.1.4, in case the matroid is ISD, becomes slightly richer; indeed, an exchange occurring on a basis also implies an exchange on the complementary basis. The next Proposition, despite straightforward, seems to be novel:

PROPOSITION 4.1.6. *Given an ISD matroid* M*, a basis* $B \in \mathcal{B}(M)$*,* $f \in B$ *and* $e \in E - B$*, the set* $(B \cup e) - f$ *is a basis of* M *if and only if* $f \in C(B, e)$ *and* $e \in C(E - B, f)$*. Moreover, every pair of circuits* C_1, C_2 *with* $|C_1 \cap C_2| = 2$ *corresponds to such an exchange.*

In [14] a few general insights are provided about ISD matroids:

PROPOSITION 4.1.7. *ISD matroids satisfy the following:*

- *The only connected, graphic ISD matroid is the trivial matroid on* $[2]$ *in which every point is a basis.*
- *If a matroid* M *on* $[2N]$ *is ISD then has at least* 2^N *bases.*

- *No contraction or restriction of a ISD matroid is an ISD matroid.*

In particular, the second statement suggests that the description of ISD matroids through the set of bases is not particularly concise. This motivates looking for more compact descriptions, of which the one presented below is an instance.

4.2. Cycles, Flats and $\mathcal{Z}$.

As mentioned, any two distinct circuits are incomparable, with respect to the partial order $\subseteq$. This suggests to consider the set of those subsets of E that can be obtained as union of a set of circuits (possibly empty):

$$\mathcal{Q} = \Big\{ \bigcup S \mid S \subseteq \mathcal{C}(M) \Big\}.$$

This sets will be called *cycles*. Clearly $\mathcal{C} \subset \mathcal{Q}$. The name makes sense in the context of *graphic matroids*, which are of no interest for the purpose of this thesis. $\mathcal{Q}$ is endowed with the natural lattice structure of the inclusion order; given $A, B \in \mathcal{Q}$:

- $A \vee B = A \cup B$
- $A \wedge B = \bigcup \{C \in \mathcal{C} : C \subseteq A \cap B\}$

The same can be carried out on the dual matroid M^*, obtaining the lattice of cocycles $\mathcal{Q}^*$ from the cocircuits $\mathcal{C}^*$. Both the lattices $\mathcal{Q}, \mathcal{Q}^*$ are subsets of the boolean algebra 2^E, so that they are mutually partially ordered by $\subseteq$. This suggests to consider the antitone Galois connection (f, f^*) defined as follows[7]:

Definition 4.2.1.

$$\begin{aligned} f : \mathcal{Q}^* &\to \mathcal{Q} \\ A &\mapsto \max_{\subseteq} \{C \in \mathcal{Q} \text{ s.t. } C \subseteq E - A\} \\ f^* : \mathcal{Q} &\to \mathcal{Q}^* \\ A &\mapsto \max_{\subseteq} \{C \in \mathcal{Q}^* \text{ s.t. } C \subseteq E - A\} \end{aligned}$$

The Galois connection (f, f^*) identifies a distinguished lattice embedded in $\mathcal{Q}$:

Definition 4.2.2. *$\mathcal{Z}(M) \subseteq \mathcal{Q}$ is the largest subset on which the map $f \circ f^*$ is the identity map. The elements of $\mathcal{Z}(M)$ are called Cyclic Flats.*

Because of the definition $\mathcal{Z}$ inherits a lattice structure from $\mathcal{Q}$; indeed, $\forall A, B \in \mathcal{Z}$:

- $A \vee B = f \circ f^*(A \cup B) \in \mathcal{Z}$
- $A \wedge B = f^* \circ f(A \cap B) \in \mathcal{Z}$

Another consequence of the definition is that $f \circ f^*$ is a closure operator; in fact, it turns out to be the restriction to $\mathcal{Q} \cup \mathcal{Q}^*$ of the closure operator cl of M. Indeed:

Proposition 4.2.3. *Given a matroid M, let $A \in \mathcal{Q}$, and suppose that $\operatorname{rank}(A) = |A| - k$, with $k \geq 1$. There exists a set of circuits $\{C_1, \dots, C_k\}$, such that $C_i \not\subset \cup_{j=1}^{i-1} C_j$ and $\cup_{j=1}^{k} C_j = A$. No larger set of circuits has the same properties.*

[7] This way of introducing the lattice of cyclic flats seems to be original. For the usual definition, see [15].

PROPOSITION 4.2.4. *Given a matroid M, let $A \in \mathcal{Q}$ and suppose that* $\mathrm{rank}_M(A) = |A| - k$. *Then $(E - A)$ is a flat of M^*, and* $\mathrm{rank}_{M^*}(E - A) = \mathrm{rank}_M(M) - k$.

In the same way $f^* \circ f$ is the restriction to $\mathcal{Q} \cup \mathcal{Q}^*$ of the kernel operator of M, $\ker(A) = \bigcup\{C \in \mathcal{C} \text{ s.t. } C \subseteq A\}$. Definition 4.2.1 can be rephrased in a few equivalent ways, such as using flats and coflats rather than cycles and cocycles (still as antitone connection), or using flats and cocycle, or cycles and coflats (monotone connection in the last cases).

Just like $\mathcal{B}$, $\mathcal{Z}$ is particularly well behaved with respect to duality. Indeed one straightforwardly obtains:

PROPOSITION 4.2.5. *The collection of cyclic flats of M^*, referred to as $\mathcal{Z}^*$, is obtained from $\mathcal{Z}$ by element-wise complementation:*

$$\mathcal{Z}^* = \{(E - F) \text{ s.t. } F \in \mathcal{Z}\} \tag{95}$$

Furthermore $\mathcal{Z}^*$ is a lattice, dual to $\mathcal{Z}$ as lattice. Indeed $\vee_{\mathcal{Z}}, \wedge_{\mathcal{Z}}$ dualize to $\wedge_{\mathcal{Z}^*}, \vee_{\mathcal{Z}^*}$ respectively.

On the other hand, while $\mathcal{B}$ provides a full description of M, $\mathcal{Z}$ doesn't, unless the rank of each cyclic flat is also known. Indeed the following result provides an independent axiomatization based on the lattice of cyclic flats decorated with their ranks.

THEOREM 4.2.6 ([15] Theorem 3.2). *Let $\mathcal{Z}$ be a collection of subsets of a set E and let* rank *be an integer valued function on $\mathcal{Z}$. There is a matroid M for which $\mathcal{Z}$ is the collection of cyclic flats and* rank *is the rank function restricted to the sets in $\mathcal{Z}$ if and only if:*

(Z0) *$\mathcal{Z}$ is a lattice under inclusion,*
(Z1) $\mathrm{rank}(0_{\mathcal{Z}}) = 0$,
(Z2) $0 < \mathrm{rank}(Y) - \mathrm{rank}(X) < |Y - X|$ *for all sets X, Y in $\mathcal{Z}$ with $X \subsetneq Y$, and*
(Z3) *for all sets X, Y in $\mathcal{Z}$,*

$$\mathrm{rank}(X) + \mathrm{rank}(Y) \geq \mathrm{rank}(X \vee Y) + \mathrm{rank}(X \wedge Y) + |(X \cap Y) - (X \wedge Y)|.$$

Notice that *(Z3)* is nothing but a restriction of the usual sub-modularity $\mathrm{rank}(X) + \mathrm{rank}(Y) \geq \mathrm{rank}(X \cup Y) + \mathrm{rank}(X \cap Y)$ to $\mathcal{Z}$. Indeed $\mathrm{rank}(X \vee Y) = \mathrm{rank}(X \cup Y)$ and, since $X \wedge Y$ is a union of circuits while $X \cap Y$ isn't, $\mathrm{rank}(X \cap Y) = \mathrm{rank}(X \wedge Y) + |(X \cap Y) - (X \wedge Y)|$.

From the same source, I report here the following further result [15, (3.2.1)], that is particularly insightful when dealing with circuits.

PROPOSITION 4.2.7. *The collection $\mathcal{C}$ of minimal subsets of E for which there is a set $X \in \mathcal{Z}$ with $C \subseteq X$ and $|C| = \mathrm{rank}(X) + 1$ is the collection of circuits of the matroid M for which $\mathcal{Z} = \mathcal{Z}(M)$.*

4.3. Cyclic Flats of ISD Matroids.

From Propositions 4.2.4, 4.2.5 and Theorem 4.2.6 one gets:

PROPOSITION 4.3.1. *A matroid M is ISD if and only if:*

(1) *$\mathcal{Z}(M)$ is closed by complementation, and*
(2) $\mathrm{rank}(E - F) = \mathrm{rank}(F) + \mathrm{rank}(M) - |F|$.

The rest of this section is somewhat a disconnected digression. I show a result that I considered on the way, but that didn't find direct application. Nevertheless it provides some insight about the structure of the lattice of cyclic flats of a ISD matroid. I show that if an ISD matroid M on $[2N]$ has no cyclic flats F such that $|F| = N$, then it can be injectively mapped onto a matroid $\not{M}$ on $[2N]$ of rank at most $\lfloor N/2 \rfloor$. Furthermore, even if there are ciclic flats F such that $|F| = N$ and $\text{rank}(F) = N - 1$, an injective map can still be constructed easily, using a pair of matroids, $\not{M}$ and a second matroid ISD on $[2N]$ that is paving.

For any cyclic flat A, let $\bar{A} = E - A$ be its complement; since A is closed by definition, there is no ambiguity with the bar notation for closure that is sometimes used.

Consider the lattice of cyclic flats $\mathcal{Z}$ of an ISD matroid M on $[2N]$. Every pair of incomparable cyclic flats $A, B \in \mathcal{Z}$ generate a sublattice containing the elements, arranged as complementary pairs:

$$\begin{array}{cc} & (\varnothing, E) \\ (A \vee B, \bar{A} \wedge \bar{B}) & (A \vee \bar{B}, \bar{A} \wedge B) \\ (A, \bar{A}) & (B, \bar{B}) \\ (A \wedge B, \bar{A} \vee \bar{B}) & (A \wedge \bar{B}, \bar{A} \vee B) \end{array} \tag{96}$$

where some of the pairs may degenerate to $(\varnothing, E)$. The lattice relationships within these cyclic flats are clear from the notation: in the non-degenerate case the corresponding Hasse diagram looks like a *rhombic dodecahedron*.

Let T denote the usual truncation of matroids. Its properties are well known:

PROPOSITION 4.3.2. [78, 7.3.10] *Let M be a matroid on E of non-zero rank and let $i \leq \text{rank}(M)$ be a non-negative integer. Then:*

$$\mathcal{I}(T^i(M)) = \{X \in \mathcal{I}(M) \mid |X| \leq r(M) - i\}.$$

Moreover the lattice of flats of $T^i(M)$ is obtained from the lattice of flats of M by removing all flats of rank exceeding $\text{rank}(M) - i - 1$, and making E the unique flat of rank $\text{rank}(M) - i$.

The following proposition establishes that for some ISD matroids one can truncate up to half of the rank and still preserve all the information:

PROPOSITION 4.3.3. *Let M be an ISD matroid on $[2N]$ such that $\mathcal{Z}(M)$ doesn't contain elements of size N. Then, consider the matroid $\not{M} := T^{\lfloor \frac{N}{2} \rfloor}(M)$. The map $M \mapsto \not{M}$ is injective.*

PROOF. The equality $\text{rank}(E - F) = \text{rank}(F) + \text{rank}(M) - |F|$ establishes that two complementary elements can only have the same rank if they both have size N. So all the pairs of complementary cyclic flats of M have different ranks.

Iterate the following operation on the set of cyclic flats of M with respect a growing index i, starting from $i = 1$:

- For every cyclic flat F s.t. $|F| = i$ eliminate the cyclic flat $E - F$; notice that $\text{rank}(F) < i$ and $\text{rank}(E - F) > \text{rank}(M) - i$.

when $i = \lfloor N/2 \rfloor$ I have eliminated only the larger element of each pair of cyclic flats $(F, E - F)$, and at the same time every cyclic flat with rank bigger than $\lceil N/2 \rceil$.

From Proposition 4.3.2 I know that the new circuits of $\not{M}$ have size larger than $\lceil N/2 \rceil$, and all flats M of size smaller than are $\lceil \text{rank}(M)/2 \rceil$ are still flats of $\not{M}$ with unchanged ranks. For every F returned by the iterative procedure above, F is still a flat of the same rank as in M, and $E - F$ is still a union of circuit as in M. Therefore F is a cyclic flat of $\not{M}$, with the same rank as in M.

□

Therefore the map $T^{\lfloor \frac{N}{2} \rfloor}$ is an embedding of the ISD matroids on $[2N]$ with no cyclic flats of size N into the matroids on $[2N]$ of rank at most $\lfloor N/2 \rfloor$.

When cyclic flats of size N are present the situation is more complicated, since the truncation may remove both cyclic flats of such pairs, depending on the rank, and some information is lost. If the cyclic flats of size N are circuit/hyperplanes this seems to be a minor problem, since circuit/hyperplanes are incomparable with any other cyclic flats besides $\varnothing, E$, therefore form disjoint sub-lattices of $\mathcal{Z}(M)$. Because of that one can think of a slightly different embedding based on the same ideas, which addresses this cases as well. On the other hand, if the cyclic flats of size N have rank smaller than $N - 1$ the last argument doesn't hold anymore, and new ideas are necessary.

One might ask how restrictive is, for a matroid M' on $[2N]$ of rank at most $\lfloor N/2 \rfloor$, the condition of being the image of an ISD matroid M via $T^{\lfloor \frac{N}{2} \rfloor}$. The condition is in fact rather restrictive, requiring a sort of *non-local submodularity*. By this I mean that the cyclic flats of M' have to satisfy *(Z3)* from [15] not only within $\mathcal{Z}(M')$, but also with respect to the *ghost* complements that have been removed by $T^{\lfloor \frac{N}{2} \rfloor}$. A complete description of this phenomenon is not ready at the moment.

REMARK 4.3.4. *It is worth noticing that the embedding of Proposition 4.3.3, despite not solving the problem completely, has the advantage of preserving representability. Indeed the truncation is a vectorial construction on* $\mathbb{C}$ [109, Proposition 7.4.10]*, therefore if* M *is representable on* $\mathbb{C}$*,* $\not{M}$ *also is.*

4.4. Representable ISD Matroids.

I already used extensively a ISD matroid that is represented by the complex matrix $[\mathbb{1}|U]$, where U is unitary. This naturally poses the question: do all representable ISD matroids exhibit this feature? The statement is proven for binary matroids [14], but, to the best of my knowledge, the same problem on complex representations is open.

In fact, concerning a representable matroid M on $[2N]$ with canonical representation $[\mathbb{1}, D]$, identically self duality corresponds to the fact that the two points on $\mathrm{G}_N(\mathbb{C}^{2N})$:

$$[\mathbb{1}, D], \qquad [D^T, \mathbb{1}]$$

belong to the same stratum. There doesn't seem to be a simple argument to tell whether a non-unitary matrix D could satisfy this property as well.

5. Abelian Matroids

In this section I define a class of ISD matroids that naturally emerges from the properties of finite abelian groups; they are highly symmetrical objects. I'm not aware of publications addressing their combinatorial properties[8]. Similar objects are studied in Code Theory, see for example [103, Chapter 11], but the latter mainly deals with linear dependence with respect to finite fields, while I will work on $\mathbb{C}$.

Within the framework outlined in the previous sections, the abelian matroids also provide a concise description of sparsity properties of graph signals, when the graph is a Cayley graph of a finite abelian group, as well as allowing the construction of multi-scale frames with localization properties.

The definition I give is based on the matrix representation, nonetheless some of the properties that I will later describe don't quite depend on the representation; indeed a distinctive feature of these matroids is that their collection of circuits includes the so called (p, q)-Dirac combs, defined below.

The main results of this section are Theorem 5.2.7, Theorem 5.2.11 and Corollary 5.2.12. The latter essentially establishes that, under one assumption, the lattice of cyclic flats of the abelian matroids is atomistic, determined by the prime factors of the size, and takes a distinguished form, that corresponds to the intersection between the union lattice generated by the atoms, and the intersection lattice generated by their complements. It's plausible that the assumption is artificial, but a more general proof is still unavailable.

One of the main motivations of my investigation was understanding if some of the properties of abelian matroids are shared by a more general class of matroids. This point is unfortunately still unclear.

5.1. Preliminaries.

DEFINITION 5.1.1. *Given a group $(G, +)$, let $S \subset G$ be a set of generators $S = \{s_1, ..., s_k\}, G =< S >$. The corresponding Cayley graph $\Gamma_{G,S}$ is is a directed graph defined as follows:*

- *Vertex set $V := G$,*
- *Edge set $E := \{(u, u + s),\ u \in V \mid s \in S\}$.*

In what follows I always assume:

ASSUMPTION 5.1.2. *G and S are such that:*

(1) *G is a finite abelian group.*
(2) *S is symmetric, that is $-S = S$, and $e \notin S$.*

If these conditions hold $(u, v) \in E$ implies $(v, u) \in E$, and I will consider them as a single undirected edge. As a result the graph is undirected, and regular of degree $d = |S|$.

EXAMPLE 5.1.3. On $G = \mathbb{Z}_m$, with $S = \{1, -1\}$, Γ is a cycle of length m. In fact whenever G is a finite abelian group and S contains a minimal set of generators and their inverses only, Γ will be a cartesian product of cycles (a torus). See for example [18].

[8] I gave the name *abelian matroid* for practical convenience. I'm not aware of other uses of this name.

DEFINITION 5.1.4. *A function $\chi : G \to \mathbb{C}^*$ is called character if it satisfies $\chi(u+v) = \chi(u)\chi(v), \forall u, v \in G$.*

The group of characters of G is denoted by $\hat{G}$. The following results are classical (see for example [26]):

PROPOSITION 5.1.5. *If $(G, +)$ is a finite abelian group, there are exactly $|G|$ characters such that they are linearly independent as vectors in $\mathbb{C}^G$. Every character satisfies $|\chi_i(v)| = 1, \forall i, v$.*

PROPOSITION 5.1.6. *Let $\{\chi_i\}_{i=1,\dots,|G|}$ be the characters of G, then:*

$$\sum_{i=1}^{|G|} \frac{\chi_i(g)}{|G|} = \mathbb{1}_e(g) \tag{97}$$

The space of graph signals $\mathcal{V}$ identifies with $\ell^2(G)$, that is $\mathbb{C}^G$ with the natural hermitian product:

$$\langle f_1, f_2 \rangle = \sum_{g \in G} f_1(g)\bar{f}_2(g).$$

The space $\ell^2(\hat{G})$ is defined in the same way.

The graph Fourier transform corresponds in this case to the group Fourier transform:

DEFINITION 5.1.7. *Let G be a finite abelian group. If $f \in \ell^2(G)$ its Fourier transform is $\hat{f} \in \ell^2(\hat{G})$:*

$$\hat{f}(\chi) = \frac{1}{\sqrt{|G|}} \sum_{g \in G} f(g)\bar{\chi}(g).$$

By (5.1.6), one has:

$$f(g) = \frac{1}{\sqrt{|G|}} \sum_{\chi \in \hat{G}} \hat{f}(\chi)\chi(g).$$

Just like in classical Fourier analysis, one can define a convolution product, which has the usual interplay with the Fourier transform:

$$(f_1 * f_2)(g) := \sum_{g'} f_1(g')f_2(g - g'), \quad (\hat{f}_1 * \hat{f}_2)(\chi) := \sum_{\chi'} \hat{f}_1(\chi')\hat{f}_2(\chi - \chi'),$$

$$\widehat{(f_1 * f_2)}(\chi) = \hat{f}_1(\chi)\hat{f}_2(\chi), \quad \widehat{(f_1 f_2)}(\chi) = (\hat{f}_1 * \hat{f}_2)(\chi).$$

The groups G and $\hat{G}$ have a natural action on $\ell^2(G)$. For $f \in \ell^2(G)$ and $g \in G, \chi \in \hat{G}$:

$$\begin{aligned} T_{\mathbb{1}_g}[f](\cdot) &= f(\cdot - g) = \sum_{g'} f(g')\mathbb{1}_g(\cdot - g') = (f * \mathbb{1}_g)(\cdot), \\ M_\chi[f](\cdot) &= \chi(\cdot)f(\cdot), \end{aligned} \tag{98}$$

corresponding to the actions on $\ell^2(\hat{G})$:

$$\begin{aligned} \widehat{T_g[f]}(\cdot) &= \widehat{(f * \mathbb{1}_g)}(\cdot) = \hat{f}(\cdot)\widehat{\mathbb{1}_g}(\cdot) =: M_{\widehat{\mathbb{1}_g}}[\hat{f}](\cdot), \\ \widehat{M_\chi[f]}(\cdot) &= \widehat{\chi f}(\cdot) = (\hat{\chi} * \hat{f})(\cdot) =: T_{\hat{\chi}}[\hat{f}](\cdot), \end{aligned} \tag{99}$$

where, for $\chi' \in \hat{G}$ I observe that $\widehat{\mathbb{1}_g}(\chi') = \bar{\chi'}(g)$, and for $g \in G$ I have $\hat{\chi}(\chi') = \mathbb{1}_\chi(\chi')$.

Roughly speaking G acts on the group domain, through T, in the same way as $\hat{G}$ acts on the frequency domain through M.

The T transform is commonly called *translation*, and M *modulation*.

PROPOSITION 5.1.8. *If H is a subgroup of a finite abelian G, for every character χ_H of H there are exactly $|G/H|$ linearly independent characters of G whose restriction to H is equal to χ_H.*

Recall the group quotient G/H is isomorphic to the set of cosets $\{g + H | g \in G\}$ equipped with the set-operation defined from the operation $+$ of G. Notice that the set of cosets $S + H \subset G/H$ generate the quotient G/H. The set of cosets $\{g + H | g \in G\}$ is formally redundant, since for $g_1, g_2 \in G$ one has that $g_1 \neq g_2$ does not imply $g_1 H \neq g_2 H$ whenever $g_2 - g_1 \in H$. It is well known that, whenever H is a normal subgroup (hence always in the abelian case), one can choose a set of representatives $\dot{G} \subset G$ with $|\dot{G}| = |G/H| = |G|/|H|$ such that the set of cosets $\{\dot{g} + H | \dot{g} \in \dot{G}\}$ is not redundant, and hence corresponds *one to one* with G/H. As a result, one can make any $g \in G$ correspond to a couple $(g_H, g_{G/H}), g_H \in H, g_{G/H} \in G/H$, having $g = g_H + \dot{g}$ with $\dot{g}$ being the representative of $g_{G/H} \equiv \dot{g} + H$. This is formalized as follows:

DEFINITION 5.1.9. *Let $\pi : G \to G/H$ be the map defined as $g \mapsto g + H$; Let $\sigma : G/H \to G$ be the map defined as $g + H \mapsto \dot{g}$.*

Notice that π is a homomorphism of groups, and it is surjective but non-injective, while σ is injective but neither surjective nor a homomorphism of groups. Also notice that $g - \sigma \circ \pi(g) \in H$, and the decomposition above reads $g \equiv (g_H, g_{G/H}) \equiv \Big(g - \sigma \circ \pi(g), \pi(g)\Big)$.

The following complements Proposition 5.1.8:

LEMMA 5.1.10. *Let a character χ_H on H be fixed, and suppose χ_G is one of its extensions as a character on G. The function $\zeta : G \to \mathbb{C}^*$ defined by:*

$$\chi_G(g) =: \chi_H(g_H)\zeta(g) \tag{100}$$

depends on $\dot{g}$ only.

PROOF.

$$\zeta(g) = \frac{\chi_G(g)}{\chi_H\Big(g - \sigma \circ \pi(g)\Big)} = \frac{\chi_G(g)}{\chi_G\Big(g - \sigma \circ \pi(g)\Big)} = \chi_G\Big(\sigma \circ \pi(g)\Big). \tag{101}$$

□

Going back to the Cayley graph of some abelian group G of cardinality n with characters $\{\chi_i\}_{i=1,\dots,n}$, notice that, by construction:

LEMMA 5.1.11. *As vectors in $\mathbb{C}^V$, every character χ_i is an eigenvector of A with eigenvalue $\alpha_i = \sum_{s \in S} \chi_i(s)$. There are no other linear independent eigenvectors of A.*

PROOF. Notice that:

$$(A\chi_i)(u) = \sum_{v \in V} A_{u,v}\chi_i(v) = \sum_{(u,v) \in E} \chi_i(v) = \sum_{s \in S} \chi_i(u + s) = \chi_i(u) \sum_{s \in S} \chi_i(s) \tag{102}$$

Exclusivity follows by counting arguments. □

Notice that α_i is real, as expected, since by construction $\chi(-v) = \chi(v)^{-1} = \overline{\chi(v)}$ (and $\chi(e) = 1$ for every character). The characters are orthogonal by symmetry of A. From Lemma 5.1.11 it follows that the two eigenvectors χ_i and $\overline{\chi_i}$ correspond to the same eigenvalue α_i, and as a result:

$$\phi_i = \chi_i + \overline{\chi_i} \tag{103}$$

is an eigenvector for A as operator on $\mathbb{R}^V$.

5.2. Definition and properties of Abelian Matroids.

DEFINITION 5.2.1. *Given a finite abelian group G with $|G| = N$, let U be a matrix having vertical indexes corresponding to elements of G, and columns corresponding to the characters of G, normalized as unit vectors of $\ell^2(G)$. The matroid $A(G)$ is the matroid on $E = [2N] \equiv G \sqcup \hat{G}$ represented by the matrix with values in $\mathbb{C}$:*

$$[\mathbb{1}|U].$$

The order in which the columns of U are presented is arbitrary but fixed.

In case G is cyclic, the matrix U takes the form of Vandermonde matrix, with respect to a root of the unity: $U_{jl} = \xi^{(j-1)(l-1)}, \xi = \exp(2\pi i/N)$; the equivalent form $U_{jl} = \xi^{jl}, \xi = \exp(2\pi i/N)$ is more practical. If G is cyclic I denote $A(G)$ by $A(N)$ with $N = |G|$.

Considering the index set $V = [N] \equiv G$, V_p denotes the set of its subsets of cardinality p such that the corresponding elements of G are in the orbit $g + H, g \in G$ of some subgroup $H < G$ with $|H| = p$. Clearly, V_p is unaffected by cyclic renaming of the indexes. Similarly W_p refers to subsets of W of the same kind, with respect to $\hat{G}$.

DEFINITION 5.2.2. *A set $F \subset E$ is called Dirac comb of type (p, q) if its intersections with V and W are in V_p and W_q respectively.*

From Proposition 3.5.5, I easily get:

COROLLARY 5.2.3. *Consider the matroid $A(N)$ with $N = pq$. Let $\{C_i\}_{i=1,\cdots,N}$ be the (p, q)-Dirac combs. Then $\{f[C_i]\}_{i=1,\cdots,N}$ is an orthonormal basis of $\mathcal{V}$.*

I denote the collection of Dirac combs of types (p, q) with the symbols $Ш(1, 0, 0, 1)$, and types (q, p) with $Ш(0, 1, 1, 0)$. It is convenient to extend this notation in the following way:

DEFINITION 5.2.4. *The symbol $Ш(i, j, k, l)$ refers to the collection of subsets of E such that, $\forall S \in Ш(i, j, k, l)$:*

- *$S \cap V$ is a union of i distinct sets in V_p and j distinct sets in V_q,*
- *$S \cap W$ is a union of k distinct sets in W_p and l distinct sets in W_q.*

In particular $Ш(0, 0, 0, 0) = \{\varnothing\}, Ш(q, 0, 0, p) = Ш(0, p, q, 0) = \{E\}$.

By construction $S \in Ш(i, 0, 0, l)$ implies $(E - S) \in Ш(q - i, 0, 0, p - l)$, and $S \in Ш(0, j, k, 0)$ implies $(E - S) \in Ш(0, p - j, q - k, 0)$. Furthermore, for any $S \in Ш(i, 0, 0, l)$ I have $|S| = ip + lq$, and for any $S \in Ш(0, j, k, 0)$ I have $|S| = jq + kp$.

PROPOSITION 5.2.5. *Every element in $Ш(1, 0, 0, 1) \cup Ш(0, 1, 1, 0)$ is a circuit. There is no other circuit of size equal to, or smaller than, $p + q$.*

PROOF. By [69, page 2] every element in $Ш(1,0,0,1) \cup Ш(0,1,1,0)$ that is supported on the union subgroups $H_V \cup H_W$ (with $H_V < G, H_W < \hat{G}$) is a circuit of $A(pq)$. Applying a translation and a modulation does not modify the size of the support. The statement follows. □

Therefore the elements of $Ш(1,0,0,1) \cup Ш(0,1,1,0)$ have rank equal to $(p+q-1)$, and any other circuit has strictly higher rank. Since $A(pq)$ is ISD, the complements of circuits are hyperplanes, therefore any element of $Ш(q-1,0,0,p-1)$ or $Ш(0,p-1,q-1,0)$ has rank $pq-1$.

Let:

$$Ш(\cdot,0,0,\cdot) := \{\varnothing, E\} \cup \bigcup_{1 \leq i < q, 1 \leq l < p} Ш(i,0,0,l),$$

and similarly for $Ш(0,\cdot,\cdot,0)$.

PROPOSITION 5.2.6. *Every element of* $Ш(\cdot,0,0,\cdot) \cup Ш(0,\cdot,\cdot,0)$ *is a cyclic flat of* $A(pq)$.

PROOF. By construction, each being union of circuits from $Ш(1,0,0,1) \cup Ш(0,1,1,0)$ and complement of a union of circuits from $Ш(1,0,0,1) \cup Ш(0,1,1,0)$ at the same time. The cases $Ш(1,0,0,0), Ш(q-1,0,0,p)$ and similar are excluded, as clearly they are independent sets, hence not cyclic flats. □

Since $Ш(\cdot,0,0,\cdot) \cup Ш(0,\cdot,\cdot,0) \subset \mathcal{Z}_{A(pq)}$ they need to have meets and joins in $\mathcal{Z}_{A(pq)}$. Clearly any pair $F_{i,l} \in Ш(i,0,0,l), F_{i',l'} \in Ш(i',0,0,l')$ yields:

$$F_{i,l} \wedge F_{i',l'} = \begin{cases} F_{i,l} \cap F_{i',l'} & \text{if } F_{i,l} \cap F_{i',l'} \in Ш(\cdot,0,0,\cdot) \\ \varnothing & \textit{otherwise.} \end{cases}$$

Despite that, $Ш(\cdot,0,0,\cdot) \cup Ш(0,\cdot,\cdot,0)$ as a set of cyclic flats is not a lattice, as the lattice operations $\wedge, \vee$ of $\mathcal{Z}_{A(pq)}$ are not yet defined for mixed pairs $F_{i,l} \in Ш(i,0,0,l), F_{j,k} \in Ш(0,j,k,0)$.

Based on these facts, I am going to describe how $Ш(\cdot,0,0,\cdot) \cup Ш(0,\cdot,\cdot,0)$, as a set of cyclic flats, can be completed into a lattice, *in a minimal way*, still satisfying all the known properties of $\mathcal{Z}_{A(pq)}$; this aims at maybe obtaining the latter, or at least a well behaved subset of it.

Mixed pairs, like $(F_{i,l} \in Ш(i,0,0,l), F_{j,k} \in Ш(0,j,k,0))$ are ordered according to:

$$(F_{i',l'}, F_{j',k'}) \subseteq (F_{i,l}, F_{j,k}) \text{ iff } F_{i',l'} \subseteq F_{i,l}, F_{j',k'} \subseteq F_{j,k}.$$

Recall that a finite lattice $\mathcal{L}$ is *atomistic*, with respect to a set of atoms $\{a_1, \ldots, a_n\} \subset \mathcal{L}$, iff every element $l \in \mathcal{L}$ can be expressed as $a_{j_1} \vee \ldots \vee a_{j_k}$ for some $\{j_1, \ldots, j_k\} \subseteq \{1, \ldots, n\}$. Also recall that, given a set of incomparable atoms $\partial\mathcal{L} \subset 2^E$, the corresponding *union lattice* is the set of subsets:

$$\mathcal{L} = \{\bigcup S \mid S \subseteq \partial\mathcal{L}\}$$

ordered by inclusion. The pointwise complementation of $\mathcal{L}$ can be obtained in the following way. Let $\overline{\partial\mathcal{L}} = \{E - A \mid A \in \partial\mathcal{L}\}$:

$$\overline{\mathcal{L}} = \{\bigcap S \mid S \subseteq \overline{\partial\mathcal{L}}\}.$$

THEOREM 5.2.7. *Let* (q,p) *be positive coprime integers. Let* M *be a ISD matroid on* $E = [2qp]$*, such that the subsets in* $Ш(1,0,0,1) \cup Ш(0,1,1,0)$ *are the only circuits of size* $q+p$*, and there is no circuit of size smaller than* $q+p$*. Then* $Ш(\cdot,0,0,\cdot) \cup Ш(0,\cdot,\cdot,0)$ *is a proper subset of* $\mathcal{Z}(M)$*;*

Consider the set of circuits $\mathcal{C}_{Ш}$ *that are obtained from the* only cyclic flats in $Ш(\cdot,0,0,\cdot) \cup Ш(0,\cdot,\cdot,0)$ *following the procedure in Proposition 4.2.7 using the ranks attained by these cyclic flats in* M*;*

If there exists a ISD matroid $M_{Ш}$ *such that,* $\forall F \in Ш(\cdot,0,0,\cdot) \cup Ш(0,\cdot,\cdot,0)$*, all the circuits of* $M_{Ш}$ *contained in* F *are in* $\mathcal{C}_{Ш}$*, then the lattice* $\mathcal{Z}(M_{Ш})$ *in obtained in the following way;*

Construct:

$$Ш_{\subseteq} = \min_{\subseteq}\Big\{(F_{i,l},F_{j,k}), F_{i,l}\in Ш(i,0,0,l), F_{j,k}\in Ш(0,j,k,0) \mid \\ ij+kl \geq \min\big(\operatorname{rank}(F_{i,l}),\operatorname{rank}(F_{j,k})\big)+1\Big\}, \tag{104}$$

where $\min_{\subseteq}$ *returns the set of incomparable minimal pairs with respect to the pair ordering. From this one gets the following set of distinguished subsets:*

$$\partial\mathcal{Z} = \{\varnothing\} \cup Ш(1,0,0,1) \cup Ш(0,1,1,0) \cup \big\{F_1\cap F_2 \mid (F_1,F_2)\in Ш_{\subseteq}\big\},$$

with $\operatorname{rank}(F_1\cap F_2) = \min\big(\operatorname{rank}(F_1),\operatorname{rank}(F_2)\big)$*; Finally,* $\mathcal{Z}(M_{Ш})$ *is the atomistic lattice:*

$$\{\cup S \mid S\subseteq \partial\mathcal{Z}\}\bigcap\{\cap S \mid S\subseteq \overline{\partial\mathcal{Z}}\},$$

with:

$$F_1\wedge F_2 = \bigcup\{F \mid F\in\partial\mathcal{Z}, F\subseteq F_1\cap F_2\},$$
$$F_1\vee F_2 = \bigcap\{F \mid (E-F)\in\partial\mathcal{Z}, F\supseteq F_1\cup F_2\},$$

for $F_1,F_2\in\mathcal{Z}$*.*

Proof. Any pair of subsets $S_1,S_2\subset V$ such that $S_1\in V_p$ and $S_2\in V_q$, satisfies $|S_1\cap S_2|=1$. In fact any element of V corresponds to such an intersection. In the same way, for any $F_1\in Ш(1,0,0,1)$ and $F_2\in Ш(0,1,1,0)$, $|F_1\cap F_2|=2$ (an exchange pair, as in Proposition 4.1.6). Hence all the elements of $\partial\mathcal{Z}$ are incomparable. Moreover:

$$|F_{i,l}\cap F_{j,k}| = ij+kl, \quad \forall F_{i,l}\in Ш(i,0,0,l), F_{j,k}\in Ш(0,j,k,0).$$

In fact $F_{i,l}\cap F_{j,k}$ has non-empty intersection with all the subsets of $F_{i,l}$ of type $Ш(1,0,0,1)$ (intersection of size $j+k$), as well as with all the subsets of $F_{j,k}$ of type $Ш(0,1,1,0)$ (intersection of size $i+l$).

On the other hand, I know from Proposition 4.2.7 that, given a cyclic flat F, each subset $A\subseteq F$ of size $|A| = \operatorname{rank}F+1$ is either a circuit or contains smaller circuits, and that all circuits are obtained in this way from some cyclic flats. Clearly any subset $A'\subseteq F$ of size $|A'|>\operatorname{rank}F+1$ must contain a circuit, and A' can be obtained as union of subsets of size $(\operatorname{rank}F+1)$, each of which either is a circuit itself or contains smaller circuits.

Therefore, whenever:

$$ij+kl \geq \min(\operatorname{rank}(F_{i,l}),\operatorname{rank}(F_{j,k}))+1 \tag{105}$$

the intersection $F_{i,l}\cap F_{j,k}$ contains some circuits (because the intersection is a subset of both cyclic flats). If equality is attained, then the intersection is itself a circuit unless it contains smaller circuits.

Let $(F_{i,l},F_{j,k})\in Ш_{\subseteq}$. It follows that, for any pair $(F_{i',l'},F_{j',k'})$ such that

$$(F_{i',l'},F_{j',k'})\subset(F_{i,l},F_{j,k}), \text{ with } F_{i',l'}\in Ш(i',0,0,l'), F_{j',k'}\in Ш(0,j',k',0)$$

one has

$$F_{i',l'} \cap F_{j',k'} \subseteq F_{i,l} \cap F_{j,k}$$

because of the ordering of pairs, and

$$i'j' + k'l' < \min(\text{rank}(F_{i',l'}), \text{rank}(F_{j',k'})) + 1$$

by construction of $\text{Ш}_{\subseteq}$. Therefore there are two possibilities:

- $F_{i,l} \cap F_{j,k}$ is a circuit, if equality is attained,
- $F_{i,l} \cap F_{j,k}$ is a union of circuits of size $ij + kl - 1$, if the strict inequality is attained.

The second statement holds because a subset of $F_{i,l} \cap F_{j,k}$ could properly contain smaller circuits only if the latter are subsets of another cyclic flat $F_{i',l'} \subseteq F_{i,l}$ or $F_{j',k'} \subseteq F_{j,k}$. But in this case one of the pairs $(F_{i',l'}, F_{j,k})$ or $(F_{i,l}, F_{j',k'})$ would be smaller than $(F_{i,l}, F_{j,k})$ in the pair ordering, and would satisfy Inequality (105), which contradicts the construction of $\text{Ш}_{\subseteq}$.

By noticing that $F_{i,l} \cap F_{j,k}$ is a flat, I have that the intersection $F_{i,l} \cap F_{j,k}$ is, in both cases above, a cyclic flat. It follows that $F_{i,l} \wedge F_{j,k} = F_{i,l} \cap F_{j,k}$, and $F_{i',l'} \wedge F_{j',k'} = \varnothing$ for any pair $(F_{i',l'}, F_{j',k'}) \subsetneq (F_{i,l}, F_{j,k})$.

On the other hand, intersections of pairs that are higher in the pair ordering can be obtained from smaller intersections thanks to the well known set theoretic identities:

$$A \cap (B \cup C) = (A \cap B) \cup (A \cup C), A \cup (B \cap C) = (A \cup B) \cap (A \cup C).$$

This way, the lattice operation $\wedge_{\mathcal{Z}}$ is now well defined on the mixed pairs. As for the lattice operation $\vee_{\mathcal{Z}}$, it is obtained from $\wedge_{\mathcal{Z}}$ by complementation.

It remains to prove that the new cyclic flats have well defined meet and join with any other cyclic flat of the completed lattice; by again using the set theoretic identities above, all these cases can be handled through the known ones, so that $\mathcal{Z}$ is indeed closed by taking meets and joins.

□

REMARK 5.2.8. *Theorem 5.2.7 explicitly mentions the fact that* $\text{Ш}(\cdot, 0, 0, \cdot)$ *and* $\text{Ш}(0, \cdot, \cdot, 0)$ *are constructed from Dirac combs of type* (p, q) *or* (q, p)*. In fact the proof only uses this fact to make sure that all elements are ciclic flats, plus that* $|F_{i,l} \cap F_{j,k}| = ij + lk$*. Therefore Theorem 5.2.7 could be formulated at a higher level of generality, for a matroid* M *admitting two subsets of* $\mathcal{Z}(M)$*, each being an atomistic lattice that is closed by taking unions, intersections and complements.*

If $\mathcal{C}_{\text{Ш}}$ *is not complete as set of circuits of a matroid, the construction of* $\partial\mathcal{Z}$ *given by Theorem 5.2.7 might fail. In that case, completing* $\text{Ш}(\cdot, 0, 0, \cdot) \cup \text{Ш}(0, \cdot, \cdot, 0)$ *into an atomistic lattice is still possible, but some of the atoms might not be intersections of elements from* $\text{Ш}(\cdot, 0, 0, \cdot) \cup \text{Ш}(0, \cdot, \cdot, 0)$.

REMARK 5.2.9. *The matroids* M *and* $M_{\text{Ш}}$ *might coincide, but Theorem 5.2.7 is not enough to conclude that. To better understand* M *and* $M_{\text{Ш}}$ *it is indeed necessary to gather more information about the ranks of the cyclic flats, as thus far we only know the ranks of the atoms and for their complements.*

I observed in Corollary 5.2.3 that the circuits in $\text{Ш}(1, 0, 0, 1) \cup \text{Ш}(0, 1, 1, 0)$ form an orthonormal basis of $\mathcal{V}$. I also recall that, in a ISD matroid, the complement of an independent set is spanning, and that within a spanning set A every independent subset $I \subseteq A$ can be completed to a basis $B \subseteq A$. This can be combined with Proposition 3.5.4 to obtain the following two results about the ranks of the cyclic flats in $\text{Ш}(\cdot, 0, 0, \cdot) \cup \text{Ш}(0, \cdot, \cdot, 0)$.

Recall that a *rectangular Vandermonde matrix* $V_{nm} = v_n^m$ with $v_{n_1} \neq v_{n_2}, \forall n_1, n_2$ has maximal rank.

THEOREM 5.2.10. *For q, p positive prime integers, let $M = A(G)$ with G cyclic, and $|G| = qp$. Then cyclic flats $F_{i,l} \in \text{Ш}(i, 0, 0, l)$ satisfy:*

$$\text{if } i = 1 \text{ or } l = 1 \text{ then } \operatorname{rank}(F_{i,l}) = |F_{i,l}| - il,$$

and for $F_{j,k} \in \text{Ш}(0, j, k, 0)$:

$$\text{if } j = 1 \text{ or } k = 1 \text{ then } \operatorname{rank}(F_{j,k}) = |F_{j,k}| - jk,$$

where $|F_{i,l}| = ip + lq$ and $|F_{j,k}| = jq + kp$.

PROOF. Given $V' \subseteq V$ and $W' \subseteq W$ let me denote by $\mathcal{V}_{V' \cup W'}$ the space of signals whose support is contained in V', and whose Fourier transform support is contained in W'.

Consider the independent set $S \in \text{Ш}(1, 0, 0, 0)$ (the same can be carried out with $S \in \text{Ш}(0, 0, 0, 1)$). There are exactly $|S|$ Dirac combs such that their support includes S, of type $\text{Ш}(1, 0, 0, 1)$, and all intersecting V on S only; let these Dirac combs be denoted by $\{C_l\}_{l=1,\dots,|S|}$. Therefore, by Corollary 5.2.3, they are an orthonormal[9] basis for the space $\mathcal{V}_{S \cup W}$, and any other Dirac comb is orthogonal to $\mathcal{V}_S$.

Let $C_1, C_2 \in \text{Ш}(1, 0, 0, 1)$ with $S = C_1 \cap C_2$. Let $Q = (C_1 \cup C_2) - S$, of course $Q \subset W$. To compute the rank of $C_1 \cup C_2$ I need to know what circuits are contained inside $C_1 \cup C_2 = S \cup Q$, which in turn correspond to signals with minimal support. Such signals belong to the space:

$$\mathcal{V}_{C_1 \cup C_2} = \mathcal{V}_{S \cup W} \cap \mathcal{V}_{V \cup Q}. \tag{106}$$

But I observed that:

$$\mathcal{V}_{S \cup W} = \operatorname{span}\{f[C_1], f[C_2]\} \oplus \operatorname{span}\{f[C_l]\}_{l \neq 1,2}.$$

Let $\{C'_l\}_l$ be the Dirac combs such that $C' \cap W \subset Q$ and $C' \cap V \not\subset S$. One has:

$$\mathcal{V}_{V \cup Q} = \operatorname{span}\{f[C_1], f[C_2]\} \oplus \operatorname{span}\{f[C'_l]\}_l.$$

Notice that the two direct sums are between orthogonal spaces; this allows to rephrase (106) as follows:

$$\mathcal{V}_{C_1 \cup C_2} = \operatorname{span}\{f[C_1], f[C_2]\} \oplus \Big(\operatorname{span}\{f[C_l]\}_{l \neq 1,2} \bigcap \operatorname{span}\{f[C'_l]\}_l \Big).$$

For the orthogonality, the second term of the direct sum is $\{0\}$, so that $\mathcal{V}_{C_1 \cup C_2} = \operatorname{span}\{f[C_1], f[C_2]\}$. Therefore I need to compute the minimal supports of linear combinations of $f[C_1], f[C_2]$.

The same reasoning applies to unions of up to $|S| - 1$ circuits $\{C_l\}_{l=1,\dots,|S|-1}$, such that $C_l \cap C_j = S, \forall l \neq j$.

In the case of $M = A(G)$ with G cyclic, the graph signals $f[C_l]$ are actually known:

$$f[C_l] = \frac{1}{\sqrt{|S|}} M_l \Big[\sum_{j \in S} \mathbb{1}_j \Big].$$

The two modulations M_l are determined by how translated $C_l \cap W$ is with respect to the corresponding subgroup of $\hat{G}$. Since $|S|$ is either q or p, $|S|$ is prime. For simplicity assume $|S| = q$.

[9] Up to one overall rescaling.

Then, for instance, $S = \{pt\}_{t=1,\dots,q}$. The $f[C_l]$ on S are complex exponentials, for $j \in V$ taking the values:

$$f[C_l](j) = \frac{1}{\sqrt{|S|}} \exp\left(2\pi i \frac{jl}{qp}\right) \sum_{j \in S} \mathbb{1}_j = \frac{1}{\sqrt{|S|}} \exp\left(2\pi i \frac{(j/p)l}{q}\right) \sum_{j \in S} \mathbb{1}_j. \tag{107}$$

Therefore, given any k signals from the collection $\{f[C_l]\}_{l=1,\dots,|S|}$, the $k \times |S|$ matrix of the corresponding values on S is a Vandermonde rectangular matrix. Since $|S|$ is prime, this matrix has maximal rank. This provides that the minimal support of a linear combination of k functions from $\{f[C_l]\}_{l=1,\dots,k}$ is $|\cup_{l=1,\dots,k} C_l| - k + 1$. Furthermore every subset of such size can be obtained, which in turn implies the claim about the rank via Proposition 4.2.3.

In the case $k = 2$ this property is referred to as C_1, C_2 being a *modular pair*, see for example [30].

This covers the cases $(i,l) = (1,l), (i,l) = (i,1), (j,k) = (1,k), (j,k) = (j,1)$.

□

Notice that the rank functions of Theorem 5.2.10 $\operatorname{rank}(F_{i,l}) = |F_{i,l}| - il$ and $\operatorname{rank}(F_{j,k}) = |F_{j,k}| - jk$ (proved so far only if one of the indexes is 1) actually satisfy both the sub-modularity and the complementarity relation $\operatorname{rank}(E - F) = \operatorname{rank}(F) + pq - |F|$, as required. Therefore they are candidate to hold for any indexes.

THEOREM 5.2.11. *For q, p distinct positive prime integers, let $M = A(G)$ with G cyclic, and $|G| = qp$. Then cyclic flats $F_{i,l} \in \text{Ш}(i,0,0,l)$ satisfy:*

$$\operatorname{rank}(F_{i,l}) = |F_{i,l}| - il,$$

and for $F_{j,k} \in \text{Ш}(0,j,k,0)$:

$$\operatorname{rank}(F_{j,k}) = |F_{j,k}| - jk,$$

where $|F_{i,l}| = ip + lq$ and $|F_{j,k}| = jq + kp$.

PROOF. The first part of the proof of Theorem 5.2.10 still holds by just choosing, for instance, $S \in \text{Ш}(i,0,0,0)$, which implies $S \subset V$ and $|S| = ip$, and $Q \in \text{Ш}(0,0,0,l)$, which implies $Q \subset W$ and $|Q| = lq$ (the same argument can start with $S \in \text{Ш}(0,0,0,l)$, or analogously in $\text{Ш}(0,\cdot,\cdot,0)$).

Therefore, for any pair (i,l) the circuits contained in F_{il} are supports of linear combinations of functions of type $\{f[C_h]\}_h$ for some Dirac combs $C_h \subset F_{il}, \forall h$. It remains to understand how small the support of these linear combination can be.

The number of distinct Dirac combs contained in F_{il} is clearly il, so this is also the number of non-zero coefficients (if some coefficients are null the obtained circuit is a subset of some smaller $F_{i'l'} \subset F_{il}$). Since $il < |S|$, by Gaussian elimination one can for sure obtain a support as small as $F_{il} - il + 1$. The claim corresponds to saying that the support cannot be any smaller than that.

One can visualize the situation with a bipartite graph, having i vertexes on the left representing the i subsets of S from $\text{Ш}(1,0,0,0)$ and l vertexes on the right representing the l subsets of Q from $\text{Ш}(0,0,0,1)$. Let $\{S_n\}_{n=1,\dots,i}, \{Q_m\}_{m=1,\dots,l}$ label these sets. Each edge of the bipartite graph represents a Dirac comb subset of F_{il}, therefore one coefficient of the linear combination. To remove h points from the support, inside some S_n (resp. Q_m), one performs the Gaussian elimination and as a result fixes some of the edges ending in S_n (resp. Q_m). Those coefficients are no longer free variables and cannot be used to perform other Gaussian eliminations at the sets Q_m on the other side of the edges (resp. S_n). *Therefore it is to prove that the values of*

the coefficients cannot be valid Gaussian elimination coefficients for both ends of one edge at the same time.

Notice that, for every S_n, the functions corresponding to the Dirac combs take, on S_n, the values in (107), therefore being a Vandermonde matrix (resp. Q_m).

Recall the following facts about Gaussian Elimination applied to Vandermonde matrices. Let V take the form:

$$V = \begin{bmatrix} 1 & v_1 & v_1^2 & \dots & v_1^{n-1} \\ 1 & v_2 & v_2^2 & \dots & v_2^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & v_n & v_n^2 & \dots & v_n^{n-1} \end{bmatrix}$$

Applying Gaussian elimination, one gets eventually the upper triangular matrix:

$$W = \begin{bmatrix} w_{1,1} & w_{1,2} & \dots & w_{1,n} \\ 0 & w_{2,2} & \dots & w_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & w_{n,n} \end{bmatrix},$$

whose h-th row is a linear combination of the first h rows of V, for any h. Moreover, for $1 \leq i \leq j \leq n$, one has:

$$w_{i,j} = M_{j-i}(\mathbf{v}_{i-1}, v_i) \prod_{k=1}^{i-1} (v_i - v_k),$$

where $\mathbf{v}_{i-1}$ is an abbreviation for $v_1, \dots, v_{i-1}$, and $M_d(\mathbf{v}_{i-1}, v_i)$ is the sum of the monomials of degree d in $\mathbf{v}_{i-1}, v_i$.

The Gaussian elimination is described step by step in the following way. Let $a_{i,j}$ be the entry of the i-th row and j-th column of this variable matrix. Before the i-th step, the entries that belong either to the i first rows or the $i-1$ first columns have the values that they will have at the end of Gaussian elimination, and, for $i \leq j \leq n$ and $i \leq h \leq n$, one has

$$a_{h,j} = M_{j-i}(\mathbf{v}_{i-1}, v_h) \prod_{k=1}^{i-1} (v_h - v_k). \tag{108}$$

This is true before the first step, and one has to prove that this remains true during Gaussian elimination. The i-th step does not change the i first rows nor the $i-1$ first columns. It changes $a_{h,i}$ to zero for $i \leq h \leq n$. For $i < j \leq n$ and $i < h \leq n$, it changes $a_{h,j}$ into $a_{h,j} - a_{i,j}a_{h,i}/a_{i,i}$. That is, the new $a_{h,j}$ is

$$\begin{aligned} a_{h,j} &= M_{j-i}(\mathbf{v}_{i-1}, v_h) \prod_{k=1}^{i-1} (v_h - v_k) - \frac{M_0(\mathbf{v}_{i-1}, v_h) \prod_{k=1}^{i-1} (v_h - v_k)}{M_0(\mathbf{v}_{i-1}, v_i) \prod_{k=1}^{i-1} (v_i - v_k)} M_{j-i}(\mathbf{v}_{i-1}, v_i) \prod_{k=1}^{i-1} (v_i - v_k) \\ &= (M_{j-i}(\mathbf{v}_{i-1}, v_h) - M_{j-i}(\mathbf{v}_{i-1}, v_i)) \prod_{k=1}^{i-1} (v_h - v_k) \\ &= M_{j-(i+1)}(\mathbf{v}_{i-1}, v_i, v_h) \prod_{k=1}^{i} (v_h - v_k), \end{aligned}$$

where it was used that, for $k > i$ one has $M_d(\mathbf{v}_{i-1}, v_h) - M_d(\mathbf{v}_{i-1}, v_i) = (v_h - v_i) M_{d-1}(\mathbf{v}_{i-1}, v_i, v_h)$. This shows that the structure of the $a_{h,j}$ as in (108) is kept during Gaussian elimination.

In particular, the overall coefficient that each row of V takes in the linear combination representing a row of W is of the form:

$$M_d(v_{h_1}, \dots) \prod_i (v_{k_i} - v_{k'_i}) \tag{109}$$

for some $d, \{h_1, \dots\}, \{k_1, \dots\}, \{k'_1, \dots\}$. This turns out to be the only necessary information for the purpose of the present proof.

The elimination can be performed towards removing points in a different order. One can check step by step that everything goes the same way, except one needs to keep track of the index i via a fixed permutation σ: for instance the assignment $v_i = \xi^i$ becomes $v_i = \xi^{\sigma(i)}$. In particular, *the overall coefficient that each row of V takes in the linear combination representing a row of W is still of the form (109) for a suitable set of indexes.*

Consider these coefficients as complex numbers in polar representation, and recall that each v_h is a power of $\exp(2\pi i/p)$ with integer exponent (if, for instance, $|S_n| = p$). Clearly $M_d(v_{h_1}, \dots)$ is also a power of $\exp(2\pi i/p)$. As for the remaining factors, one has:

$$\exp\left(2\pi i \frac{j}{p}\right) - \exp\left(2\pi i \frac{k}{p}\right) = 2 \exp\left(2\pi i \frac{k+j}{2p}\right) \cos\left(2\pi \frac{j-k}{2p}\right)$$

so that the overall phase factor of every coefficient is a power of $\exp(2\pi i \frac{1}{2p})$.

On the other hand, $|S_n| = p$ implies $|Q_m| = q$, and the same coefficients therein must have phase factor that is a power of $\exp(2\pi i \frac{1}{2q})$.

Clearly two complex numbers can only be equal if they have the same phase.

Since p, q are prime, this allows to conclude that coefficients of Gauss eliminations, performed on some S_n and some Q_m, can not have consistent complex phase with respect to both p and q, unless the phase is 0 (i.e. real).

But considering the form of the v_h, in a set of coefficients of a Gauss elimination performed on a set in $Ш(1,0,0,0)$ (resp.$Ш(0,0,0,1)$), there can be just one that is real. Therefore, going back to the bipartite graph described above, each edge/coefficient is only valid as Gauss elimination coefficient on one of its ends.

This proves that one can remove exactly $il - 1$ points form the support, using the il coefficients, and that in fact any set of $il - 1$ points can be removed, as long as the resulting support is not a subset of a smaller $F_{i'l'} \subset F_{il}$. The claim follows.

□

Finally I can conclude the follwing:

COROLLARY 5.2.12. *Under the hypotheses of Theorem 5.2.11, and referring to the notation of Theorem 5.2.7, the collection of circuits $\mathcal{C}_{Ш}$ is the entire collection of circuits of $A(qp)$; in other words, $M_{Ш} = M = A(qp)$ in this case.*

PROOF. The proof of Theorem 5.2.11 shows more than its claim. Indeed, it shows that all the circuits contained in a cyclic flat $F \in Ш(\cdot, 0, 0, \cdot) \cup Ш(0, \cdot, \cdot, 0)$ have the same size unless

they are contained in a smaller cyclic flat from $Ш(\cdot, 0, 0, \cdot) \cup Ш(0, \cdot, \cdot, 0)$, and all subsets of F of size $\operatorname{rank}(F) + 1$ are circuits, unless they are contained in a smaller cyclic flat from $Ш(\cdot, 0, 0, \cdot) \cup Ш(0, \cdot, \cdot, 0)$. This yields the claim, via Proposition 4.2.7. □

Corollary 5.2.12 opens a number of questions. In particular:

- For a general positive integer N, $\mathcal{Z}(A(N))$ contains as subset $\mathcal{Z}(A(pq))$ for all pairs of distinct prime factors of N, suggesting the possibility of a structure theorem for the abelian matroids. Despite that the proposed proofs do not address the case $p = q$ or the case of more than two prime factors.
- It is still unclear what are the consequences, for a general matroid, of having a lattice of cyclic flats with the strong structure that was proven for $A(pq)$ for p, q distinct and prime.

6. Spectra

In Section 2 I informally observed how some aspects of Graph Signal Processing only depend on the eigenfunctions of the graph laplacian operator $\mathcal{L}$, while others on the contrary depend on the eigenvalues as well. In the previous sections the focus has fully been on a set of eigenfunctions $\{u_j\}_{j\in W}$. This last section is an attempt to address the combinatorial relationship between eigenfunctions and eigenvalues (if any). The framework seems to be unprecedented; I don't start from a graph but from a matrix U. I offer a few results and, lastly, an open problem that requires more research.

I already commented on the fact that the laplacian eigenvalues are called frequencies because of the formal analogy with time-frequency analysis. Indeed, using the definition, I have:

$$\lambda_j = \langle u_j, \mathcal{L}u_j\rangle = \langle u_j, (D-A)u_j\rangle = \sum_{(l,k)} A_{lk}|u_j(l) - u_j(k)|^2 \tag{110}$$

where the last sum is over the ordered pairs of indexes. This shows that 0 is always an eigenvalue, whose eigenfunction(s) is(are) constant on connected components. This formula also allows to observe that each eigenvalue can be regarded as an estimate of how sharply the corresponding eigenfunction(s) varies on the graph, motivating the name *frequency* referred to an eigenvalue.

6.1. Non-Trivial Spectrum Existence.

The vast majority of the available research about the Graph Laplacian focuses on the spectrum and its properties, once the underlaying graph is fixed and known. In this thesis I take a turn on this framework, by focusing primarily on the eigenfunctions. In fact, given U, neither the graph nor the spectrum are determined, and this section elaborates on the set of graphs (therefore spectra) compatible with U.

I now consider a given matrix $U \in U(N)$ and investigate which spectra can arise from it. Columns and rows of U will be labelled by indices in V and W respectively. A spectrum will be defined as a vector with positive entries presented as a diagonal matrix $\Lambda \in (\mathbb{R}_0^+)^W \hookrightarrow M_{W\times W}(\mathbb{R})$ (it will be clear whether Λ is considered a matrix or a vector case by case). Constraints are given by the properties that the Graph Laplacian $\mathcal{L}$ must satisfy:

Definition 6.1.1. *The matrices of the following set:*

$$\mathbb{L} = \{\mathcal{L} \in M_{V\times V}(\mathbb{R}) \,|\, \mathcal{L}_{ij} \leq 0, \mathcal{L}_{ii} = -\sum_{k\neq i} \mathcal{L}_{ik}\, \forall i,j \in V, i \neq j\}$$

will be called admissible laplacian operators on V.

I am imposing that $U\Lambda U^*$ is an admissible laplacian operator. The null spectrum always trivially satisfies the constraints. The following statement provides necessary and sufficient conditions for a non-trivial spectrum to exist. A partition of the set of vertexes $\{V_i\}_{i=1,\dots,k}$ is such that $V_i \subseteq V, V_i \cap V_j = \varnothing\, \forall i \neq j$ and $\cup_i V_i = V$. Partitions are partially ordered by coarseness.

Proposition 6.1.2. *Given a unitary matrix $U \in U(N)$, let $\{V_i\}_{i=1,\dots,k}$ be a maximally coarse partition of V such that k columns of U are piece-wise constant on $\{V_i\}_{i=1,\dots,k}$. If $k < N$ then there exists at least one non-trivial spectrum $\Lambda = \mathrm{diag}\{\lambda_i \geq 0\}_{i\in V}$ such that $U\Lambda U^*$ is an admissible laplacian operator, and has $(N-k)$ non-zero eigenvalues. In fact any spectrum yielding an admissible laplacian operator $U\Lambda U^*$ also has k eigenvalues equal to zero.*

PROOF. Let $k = 1$, and, without loss of generality, let u_1 be the column that is constant on V. The definition of unitary matrix reads:

$$l \neq j \implies \Re\left(\sum_i u_i(l)\bar{u}_i(j)\right) = 0 = \Re(u_1(l)\bar{u}_1(j)) + \sum_{i\neq 1}\Re(u_i(l)\bar{u}_i(j))$$

since u_1 is constant, and is a column of a unitary matrix, it follows that $\Re(u_1(l)\bar{u}_1(j)) = 1/N$ and hence:

$$\sum_{i\neq 1}\Re(u_i(l)\bar{u}_i(j)) = -1/N.$$

Now let $\lambda_1 = 0$ and $\lambda_i = 1\,\forall i \neq 1$. This implies:

$$\mathcal{L}_{lj} = \sum_i u_i(l)\lambda_i\bar{u}_i(j) = \sum_{i\neq 1}\Re(u_i(l)\bar{u}_i(j)) < 0.$$

It remains to show that $\mathcal{L}_{jj} = -\sum_{i\neq j}\mathcal{L}_{ij}$. Again by orthogonality:

$$\sum_i |u_i(j)|^2 = 1 = |u_1(j)|^2 + \sum_{i\neq 1}|u_i(j)|^2$$

and $|u_1(j)|^2 = 1/N$ so that $\mathcal{L}_{jj} = \sum_i |u_i(j)|^2\lambda_i = \sum_{i\neq 1}|u_i(j)|^2 = 1 - 1/N = (N-1)(1/N)$. Therefore the admissibility conditions hold.

Let $1 < k < N$, and, without loss of generality, let $u_1, \ldots, u_k$ be the eigenfunctions that are piece-wise constant on V. Being $u_1, \ldots, u_k$ orthonormal and piece-wise constant, for $l, j \leq k$ one has:

$$\delta_{lj} = \langle u_l, u_j\rangle = \sum_{t=1}^{N} u_l(t)\bar{u}_j(t) = \sum_{i=1}^{k}|V_i|u_l(t_i)\bar{u}_j(t_i) \tag{111}$$

where $t_i \in V_i$ is an arbitrary index in the corresponding cell. Therefore the set of k vectors in $\mathbb{C}^k$ with components $\tilde{u}_j(i) = \sqrt{V_i}u_j(t_i)$ is orthonormal. As a result, the (111) holds for its columns as well:

$$\delta_{lj} = \sum_{i=1}^{k}|V_i|u_l(t_i)\bar{u}_j(t_i) = \sum_{i=1}^{k}\sqrt{|V_l||V_j|}u_i(t_l)\bar{u}_i(t_j) \tag{112}$$

Then, for any j, l:

$$\Re\left(\sum_n u_n(l)\bar{u}_n(j)\right) = 0 = \sum_{n=1}^{k}\Re(u_n(l)\bar{u}_n(j)) + \sum_{n=k+1}^{N}\Re(u_n(l)\bar{u}_n(j)),$$

let $p(j)$ be the index of the cell containing j. Then by carefully applying (112):

$$\begin{aligned}\sum_{n=1}^{k}\Re(u_n(l)\bar{u}_n(j)) = \frac{1}{\sqrt{|V_{p(l)}||V_{p(j)}|}}\sum_{n=1}^{k}\sqrt{|V_{p(l)}||V_{p(j)}|}\Re(u_n(t_{p(l)})\bar{u}_n(t_{p(j)})) = \\ \frac{1}{\sqrt{|V_{p(l)}||V_{p(j)}|}}\sum_{n=1}^{k}|V_n|\Re(\bar{u}_{p(l)}(t_n)u_{p(j)}(t_n)) = \frac{\delta_{p(l)p(j)}}{\sqrt{|V_{p(l)}||V_{p(j)}|}} \geq 0\end{aligned} \tag{113}$$

Now let $\lambda_i = 0$ for $i = 1, \dots, k$ and $\lambda_i = 1 \, \forall i > k$. This implies, for $j, l \leq k$:

$$\mathcal{L}_{lj} = \sum_n u_n(l)\lambda_n \bar{u}_n(j) = \sum_{n=k+1}^{N} \Re(u_n(l)\bar{u}_n(j)) \leq 0.$$

It remains to show that $\mathcal{L}_{jj} = -\sum_{i \neq j} \mathcal{L}_{ij}$. Again by orthogonality:

$$\sum_n |u_n(j)|^2 = 1 = \sum_{n=1}^{k} |u_n(j)|^2 + \sum_{n=k+1}^{N} |u_n(j)|^2$$

and

$$\sum_{n=1}^{k} |u_n(j)|^2 = \frac{1}{|V_{p(j)}|} \sum_{n=1}^{k} |V_{p(j)}||u_n(t_{p(j)})|^2 = \frac{1}{|V_{p(j)}|} \sum_{n=1}^{k} |V_n||u_{p(j)}(t_n)|^2 = \frac{1}{|V_j|},$$

so that $\mathcal{L}_{jj} = \sum_i |u_i(j)|^2 \lambda_i = \sum_{n=k+1}^{N} |u_n(j)|^2 = 1 - \frac{1}{|V_j|} = (|V_j| - 1)(1/|V_j|)$. Considering the block structure emerged in (113) the admissibility conditions hold.

The last statement is due to Equation (110) for $\lambda = 0$.

□

The last result showed that it suffices to constrain the matrix U to have a constant vector, and this guarantees that the corresponding graphs are non-trivial. This, roughly speaking, means working with $U(N-1)$. In what follows, for simplicity, I just consider a generic matrix in $U(N)$.

6.2. Stratification of $U(N)$ through Spectra.

Besides the existence of one spectrum, as in Proposition 6.1.2, I am now looking for a more detailed description of the whole set of spectra that provide an admissible Laplacian operator, with respect to a given $U \in U(N)$, therefore corresponding to graphs. The content of this last part of the thesis is fragmentary and conjectural, therefore I chose to present it briefly without introducing too much new notation, as this would not be justified by matching technical results.

Every spectrum yields $\mathcal{L}(\Lambda) = U\Lambda U^* \in M_{V \times V}(\mathbb{R})$ that determines A, hence a graph with vertexes labelled by V. Equality of graphs will occur when labels and weights of all the edges coincide (I do not consider equal graphs being only isomorphic).

DEFINITION 6.2.1. *Two graphs, that is, two adjacency matrixes, will be called equivalent if they turn equal as soon as every nonzero weight is replaced by* 1. *In other words, equivalent graphs have the same sparsity pattern of* $\mathcal{L}$.

From Definition 6.1.1, I see that each of the $N(N-1)/2$ admissibility condition on off-diagonal elements of $\mathcal{L}$ is an inequality ($\leq$), linear with respect to the spectrum (and quadratic with respect to U entries). The diagonal constraints are linear equations, that are satisfied whenever U fulfils the hypotheses of Proposition 6.1.2. The locus of admissible spectra is a convex subset of $(\mathbb{R}_0^+)^W$ delimited by hyperplanes through the origin, i.e. a cone. Under the hypotheses of this proposition, this cone is non-empty. Facets of this cone will have a certain subset of the off-diagonal constraints *active* (i.e. satisfied by equality), hence will correspond to a class of equivalent graphs. If a facet has dimension $N - k$ corresponds to spectra such that (at least) k

constraints are active (intersection of a facet by a hyperplane can only decrease its dimension by one at most). An active constraint corresponds to a missing edge among the possible $N(N-1)/2$ of the graph described by $\mathcal{L}$, so that 1-facets of the cone correspond to (locally) maximal sparsity of the corresponding graphs.

In this way I assigned to every $U \in U(N)$ a cone $\Delta(U) \subset (\mathbb{R}_0^+)^W$.

DEFINITION 6.2.2. *I define an equivalence relation among the cones* $\{\Delta(U) \mid U \in U(N)\}$ *by saying that* $\Delta \sim \Delta'$ *if for every facet of* Δ *there is a facet of* Δ' *with the same active constraints. The equivalence classes of* $\sim$ *divide* $U(N)$ *in a discrete set of regions, that I call strata in analogy to the ones corresponding to the matroids* $[\mathbb{1}|U]$.

REMARK 6.2.3. *The example of Abelian Matroids provides some insights:*

- *Eigenfunctions of the graph laplacian on any Cayley graph from a finite abelian group* G *and a subset* $H \subset G$ *do not depend on* H.
- *Different choices of* H *correspond to different spectra corresponding to the same group* G *hence the same* U. *Therefore the corresponding spectra are in the same cone. Depending on the sparsity induced by* H *they may correspond to edges of the cone* $\Delta(U)$.
- *Every element of* H *corresponds to a subgroup of* G.
- *Because of this I may expect a* $\Delta(U)$ *stratum to correspond to a* $M(U)$ *stratum, as subgroup information is encoded, on the one hand, in facets of* $\Delta(U)$, *on the other, in the cyclic flats of* $M(U)$.

In particular, the last point rises the question about a plausible strong relation between the two stratifications mentioned, having one finer than the other, or having both equivalent.

This motivates looking for a convenient representation of the stratification provided by spectra.

Let $E = \binom{V}{2}$, whose elements will be denoted as unordered couples of indices $(i,j), i,j \in V$. Let $O(U) \in M_{W \times E}(\mathbb{R})$ be a real $N \times \binom{N}{2}$ real matrix with entries $O_{k,(i,j)}(U) = \Re(u_k(i)\bar{u}_k(j))$. According to this notation, the off-diagonal constraints for Λ become, for each couple (i,j):

$$\sum_k \Lambda_k O_{k,(i,j)}(U) \leq 0.$$

Using the language of *oriented matroids*, the the matrix $O(U)$ represents an oriented matroid, and the facets of the cone $\Delta(U)$ correspond to its *covectors* having entries in $\{-,0\}$. Therefore an equivalence class in $U(N)/\sim$ corresponds to a class of oriented matroids on E whose subset of $\{-,0\}$-supported covectors are the same, and describe $\Delta(U)$; see for example [112].

REMARK 6.2.4. *A strong relationship between the* $\Delta(U)$ *stratification and the* $[\mathbb{1}|U]$ *stratification, if confirmed, might have a remarkable practical relevance. Indeed the former is partly related to the graph structure directly* as is, *possibly allowing to bypass the brute-forse computation of* U *and of it's matroidal counterpart.*

Bibliography

[1] Ameya Agaskar and Yue M Lu. A spectral graph uncertainty principle. *IEEE Transactions on Information Theory*, 59(7):4338–4356, 2013.

[2] Michal Aharon, Michael Elad, and Alfred Bruckstein. -svd: An algorithm for designing overcomplete dictionaries for sparse representation. *Signal Processing, IEEE Transactions on*, 54(11):4311–4322, 2006.

[3] William K Allard, Guangliang Chen, and Mauro Maggioni. Multi-scale geometric methods for data sets II: Geometric multi-resolutione analysis. *Appl. Comput. Harmon. Anal.*, 32(3):435–462, 2012.

[4] Aamir Anis, Akshay Gadde, and Antonio Ortega. Towards a sampling theorem for signals on arbitrary graphs. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3864–3868. IEEE, 2014.

[5] David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035. Society for Industrial and Applied Mathematics, 2007.

[6] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.

[7] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems*, volume 14, pages 585–591, 2001.

[8] Richard Ernest Bellman and Howard A Osborn. Dynamic programming and the variation of green's functions. 1957.

[9] Yoshua Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.

[10] Adi Benisrael. Projectors on intersections of subspaces. *Contemporary Mathematics*, pages 41–50, 2015.

[11] Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pages 97–104. ACM, 2006.

[12] Peter Binev, Albert Cohen, Wolfgang Dahmen, Ronald DeVore, and Vladimir Temlyakov. Universal algorithms for learning theory part i: piecewise constant functions. *Journal of Machine Learning Research*, 6(Sep):1297–1321, 2005.

[13] Gilles Blanchard, Olivier Bousquet, and Laurent Zwald. Statistical properties of kernel principal component analysis. *Machine Learning*, 66(2-3):259–294, 2007.

[14] John A Bondy and Dominic JA Welsh. Some results on transversal matroids and constructions for identically self-dual matroids. *The Quarterly Journal of Mathematics*, 22(3):435–451, 1971.

[15] Joseph E Bonin and Anna De Mier. The lattice of cyclic flats of a matroid. *Annals of combinatorics*, 12(2):155–170, 2008.

[16] Alexandre V Borovik, Izrail? Moiseevich Gel?fand, and Neil White. Coxeter matroids. In *Coxeter Matroids*, pages 151–197. Springer, 2003.

[17] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013.

[18] Andries E Brouwer and Willem H Haemers. *Spectra of graphs*. Springer Science & Business Media, 2011.

[19] Guillermo Canas, Tomaso Poggio, and Lorenzo Rosasco. Learning manifolds with k-means and k-flats. In *Advances in Neural Information Processing Systems*, pages 2465–2473, 2012.

[20] Guillermo Canas and Lorenzo Rosasco. Learning probability measures with respect to optimal transport metrics. In *Advances in Neural Information Processing Systems*, pages 2492–2500, 2012.

[21] Emmanuel J Candes and Justin Romberg. Quantitative robust uncertainty principles and optimally sparse decompositions. *Foundations of Computational Mathematics*, 6(2):227–254, 2006.

[22] Siheng Chen, Rohan Varma, Aliaksei Sandryhaila, and Jelena Kovačević. Discrete signal processing on graphs: Sampling theory¡? pub _newline=""? *IEEE transactions on signal processing*, 63(24):6510–6523, 2015.

[23] Michael Christ. A $T(b)$ theorem with remarks on analytic capacity and the Cauchy integral. *Colloq. Math.*, 60/61(2):601–628, 1990.
[24] Albert Cohen, Wolfgang Dahmen, Ingrid Daubechies, and Ronald DeVore. Tree approximation and optimal encoding. *Applied and Computational Harmonic Analysis*, 11(2):192–226, 2001.
[25] Ronald R Coifman and Mauro Maggioni. Diffusion wavelets. *Applied and Computational Harmonic Analysis*, 21(1):53–94, 2006.
[26] Keith Conrad. Characters of finite abelian groups. *Lecture Notes*, 2010.
[27] Thierry Coulhon, Gerard Kerkyacharian, and Pencho Petrushev. Heat kernel generated frames in the setting of dirichlet spaces. *Journal of Fourier Analysis and Applications*, 18(5):995–1066, 2012.
[28] Ingrid Daubechies. *Ten lectures on wavelets*, volume 61. Siam, 1992.
[29] Ingrid Daubechies. The work of yves meyer. In *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures*, pages 114–124. World Scientific, 2010.
[30] Emanuele Delucchi. Modular elimination in matroids and oriented matroids. *European Journal of Combinatorics*, 32(3):339–343, 2011.
[31] David L Donoho. High-dimensional data analysis: The curses and blessings of dimensionality. *AMS Math Challenges Lecture*, pages 1–32, 2000.
[32] David L. Donoho and Carrie Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proc. Natl. Acad. Sci. USA*, 100(10):5591–5596, 2003.
[33] David L Donoho and Philip B Stark. Uncertainty principles and signal recovery. *SIAM Journal on Applied Mathematics*, 49(3):906–931, 1989.
[34] Roger Duke. Matroid erection and duality. *European Journal of Combinatorics*, 8(4):367–370, 1987.
[35] Venkatesan N Ekambaram, Giulia Fanti, Babak Ayazifar, and Kannan Ramchandran. Critically-sampled perfect-reconstruction spline-wavelet filterbanks for graph signals. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 475–478. IEEE, 2013.
[36] Michael Elad and Alfred M Bruckstein. A generalized uncertainty principle and sparse representation in pairs of bases. *IEEE Transactions on Information Theory*, 48(9):2558–2567, 2002.
[37] Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2004.
[38] Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *arXiv preprint arXiv:1310.0425*, 2013.
[39] Gerald B Folland and Alladi Sitaram. The uncertainty principle: a mathematical survey. *Journal of Fourier analysis and applications*, 3(3):207–238, 1997.
[40] Israel M Gelfand, R Mark Goresky, Robert D MacPherson, and Vera V Serganova. Combinatorial geometries, convex polyhedra, and schubert cells. *Advances in Mathematics*, 63(3):301–316, 1987.
[41] Allen Gersho and Robert M Gray. *Vector quantization and signal compression*, volume 159. Springer Science & Business Media, 2012.
[42] Giacomo Gigante and Paul Leopardi. Diameter bounded equal measure partitions of Ahlfors regular metric measure spaces. *Discrete Comput. Geom.*, 57(2):419–430, 2017.
[43] Franziska Göbel, Gilles Blanchard, and Ulrike von Luxburg. Construction of tight frames on graphs and application to denoising. *arXiv preprint arXiv:1408.4012*, 2014.
[44] Alexander N Gorban. *Principal manifolds for data visualization and dimension reduction*, volume 58. Springer, 2007.
[45] Rémi Gribonval and Morten Nielsen. Sparse representations in unions of bases. *IEEE transactions on Information theory*, 49(12):3320–3325, 2003.
[46] Misha Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. Birkhäuser Boston, Inc., Boston, MA, english edition, 2007.
[47] Peter M Gruber. Optimum quantization and its applications. *Adv. Math.*, 186(2):456–497, 2004.
[48] David K Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011.
[49] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer New York Inc., 2001.
[50] Matthias Hein, Jean-Yves Audibert, and Ulrike Von Luxburg. Graph laplacians and their convergence on random neighborhood graphs. *arXiv preprint math/0608522*, 2006.
[51] Lars Hörmander. *The analysis of linear partial differential operators I: Distribution theory and Fourier analysis*. Springer, 2015.

[52] Michael E Houle, Hans-Peter Kriegel, Peer Kröger, Erich Schubert, and Arthur Zimek. Can shared-neighbor distances defeat the curse of dimensionality? In *Scientific and Statistical Database Management*, pages 482–500. Springer Berlin Heidelberg, 2010.

[53] Jeff Irion and Naoki Saito. Applied and computational harmonic analysis on graphs and networks. In *Wavelets and Sparsity XVI*, volume 9597, page 95971F. International Society for Optics and Photonics, 2015.

[54] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An introduction to statistical learning*, volume 112. Springer, 2013.

[55] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.

[56] Kenneth Kreutz-Delgado, Joseph F Murray, Bhaskar D Rao, Kjersti Engan, Te-Won Lee, and Terrence J Sejnowski. Dictionary learning algorithms for sparse representation. *Neural computation*, 15(2):349–396, 2003.

[57] Stephane Lafon and Ann B Lee. Diffusion maps and coarse-graining: A unified framework for dimensionality reduction, graph partitioning, and data set parameterization. *IEEE transactions on pattern analysis and machine intelligence*, 28(9):1393–1403, 2006.

[58] Jeffrey C Lagarias. Mathematical quasicrystals and the problem of diffraction. *Directions in mathematical quasicrystals*, 13:61–93, 2000.

[59] Wenjing Liao and Mauro Maggioni. Adaptive geometric multiscale approximations for intrinsically low-dimensional data. *J. Mach. Learn. Res.*, 20(98):1–83, 2019.

[60] Anna V Little, Mauro Maggioni, and Lorenzo Rosasco. Multiscale geometric methods for data sets i: Multiscale svd, noise and curvature. 2012.

[61] Pengfei Liu, Xiaohan Wang, and Yuantao Gu. Coarsening graph signal with spectral invariance. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1070–1074. IEEE, 2014.

[62] Stuart Lloyd. Least squares quantization in PCM. *IEEE Trans. Inform. Theory*, 28(2):129–137, 1982.

[63] Mauro Maggioni, Stanislav Minsker, and Nate Strawn. Dictionary learning and non-asymptotic bounds for geometric multi-resolution analysis. *PAMM. Proc. Appl. Math. Mech.*, 14(1):1013–1016, 2014.

[64] Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999.

[65] Stephane G Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE transactions on pattern analysis and machine intelligence*, 11(7):674–693, 1989.

[66] Andreas Maurer. Concentration inequalities for functions of independent variables. *Random Structures Algorithms*, 29(2):121–138, 2006.

[67] Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample-variance penalization. In *COLT*, 2009.

[68] Dillon Mayhew, Mike Newman, and Geoff Whittle. Yes, the ?missing axiom? of matroid theory is lost forever. *Transactions of the American Mathematical Society*, 370(8):5907–5929, 2018.

[69] Roy Meshulam. An uncertainty inequality for finite abelian groups. *European Journal of Combinatorics*, 27(1):63–67, 2006.

[70] Boaz Nadler, Stéphane Lafon, Ronald R Coifman, and Ioannis G Kevrekidis. Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. *Appl. Comput. Harmon. Anal.*, 21(1):113–127, 2006.

[71] Yuji Nakatsukasa, Naoki Saito, and Ernest Woei. Mysteries around the graph laplacian eigenvalue 4. *Linear Algebra and its Applications*, 438(8):3231–3246, 2013.

[72] Sunil K Narang and Antonio Ortega. Compact support biorthogonal wavelet filterbanks for arbitrary undirected graphs. *IEEE transactions on signal processing*, 61(19):4673–4685, 2013.

[73] Hariharan Narayanan and Sanjoy Mitter. Sample complexity of testing the manifold hypothesis. *Advances in Neural Information Processing Systems*, pages 1786–1794, 2010.

[74] Mark Newman. *Networks*. Oxford university press, 2018.

[75] Ha Q Nguyen and Minh N Do. Downsampling of signals on graphs via maximum spanning trees. *IEEE Transactions on Signal Processing*, 63(1):182–191, 2014.

[76] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.

[77] Bruno Olshausen and David Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(607):6583, 1996.

[78] James G Oxley. *Matroid theory*, volume 3. Oxford University Press, USA, 2006.

[79] Nathanael Perraudin, Benjamin Ricaud, David Shuman, and Pierre Vandergheynst. Global and local uncertainty principles for signals on graphs. *arXiv preprint arXiv:1603.03030*, 2016.

[80] Isaac Pesenson. Sampling in paley-wiener spaces on combinatorial graphs. *Transactions of the American Mathematical Society*, 360(10):5603–5627, 2008.

[81] Isaac Z Pesenson and Meyer Z Pesenson. Sampling, filtering and sparse approximations on combinatorial graphs. *Journal of Fourier Analysis and Applications*, 16(6):921–942, 2010.

[82] Peter Petersen. *Riemannian geometry*. Springer, Cham, third edition, 2016.

[83] I. F. Pinelis and A. I. Sakhanenko. Remarks on inequalities for probabilities of large deviations. *Theory Probab. Appl.*, 30(1):143–148, 1985.

[84] Iosif Pinelis. Optimum bounds for the distributions of martingales in Banach spaces. *Ann. Probab.*, 22(4):1679–1706, 1994.

[85] Iosif Pinelis. Correction: "Optimum bounds for the distributions of martingales in Banach spaces" [Ann. Probab. **22** (1994), no. 4, 1679–1706; MR1331198 (96b:60010)]. *Ann. Probab.*, 27(4):2119, 1999.

[86] Gilles Puy, Nicolas Tremblay, Rémi Gribonval, and Pierre Vandergheynst. Random sampling of bandlimited signals on graphs. *Applied and Computational Harmonic Analysis*, 2016.

[87] Idan Ram, Michael Elad, and Israel Cohen. Generalized tree-based wavelet transform. *IEEE Transactions on Signal Processing*, 59(9):4199–4209, 2011.

[88] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

[89] Aliaksei Sandryhaila and Jose MF Moura. Big data analysis with signal processing on graphs: Representation and processing of massive data sets with irregular structure. *IEEE Signal Processing Magazine*, 31(5):80–90, 2014.

[90] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.

[91] Bernhard Schölkopf, Alexander J. Smola, and Klaus-Robert Müller. Advances in kernel methods. In Bernhard Schölkopf, Christopher J. C. Burges, and Alexander J. Smola, editors, *International conference on artificial neural networks*, chapter Kernel Principal Component Analysis, pages 327–352. MIT Press, Cambridge, MA, USA, 1999.

[92] Laurent Schwartz. *Cours d'analyse. 3*. Hermann, Paris, second edition, 1993.

[93] Clayton Scott, Robert D Nowak, et al. Minimax-optimal classification with dyadic decision trees. *IEEE Trans. Inform. Theory*, 52(4):1335–1353, 2006.

[94] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.

[95] David I Shuman, Mohammad Javad Faraji, and Pierre Vandergheynst. A multiscale pyramid transform for graph signals. *IEEE Transactions on Signal Processing*, 64(8):2119–2134, 2015.

[96] David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE signal processing magazine*, 30(3):83–98, 2013.

[97] David I Shuman, Benjamin Ricaud, and Pierre Vandergheynst. Vertex-frequency analysis on graphs. *Applied and Computational Harmonic Analysis*, 40(2):260–291, 2016.

[98] David I Shuman, Christoph Wiesmeyr, Nicki Holighaus, and Pierre Vandergheynst. Spectrum-adapted tight graph wavelet and vertex-frequency frames. *IEEE Transactions on Signal Processing*, 63(16):4223–4235, 2015.

[99] Arthur Szlam. Asymptotic regularity of subdivisions of euclidean domains by iterated pca and iterated 2-means. *Applied and Computational Harmonic Analysis*, 27(3):342–350, 2009.

[100] Terence Tao. An uncertainty principle for cyclic groups of prime order. *arXiv preprint math/0308286*, 2003.

[101] Oguzhan Teke and Palghat P Vaidyanathan. Uncertainty principles and sparse eigenvectors of graphs. *IEEE Transactions on Signal Processing*, 65(20):5406–5420, 2017.

[102] Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[103] Audrey Terras. *Fourier analysis on finite groups and applications*. Number 43. Cambridge University Press, 1999.

[104] Nicolas Tremblay, Pierre-Olivier Amblard, and Simon Barthelmé. Graph sampling with determinantal processes. In *2017 25th European Signal Processing Conference (EUSIPCO)*, pages 1674–1678. IEEE, 2017.

[105] Nicolás García Trillos, Moritz Gerlach, Matthias Hein, and Dejan Slepčev. Error estimates for spectral convergence of the graph laplacian on random geometric graphs toward the laplace–beltrami operator. *Foundations of Computational Mathematics*, pages 1–61, 2019.

[106] Mikhail Tsitsvero, Sergio Barbarossa, and Paolo Di Lorenzo. Signals on graphs: Uncertainty principle and sampling. *IEEE Transactions on Signal Processing*, 64(18):4845–4860, 2016.

[107] Michael Unser. Sampling-50 years after shannon. *Proceedings of the IEEE*, 88(4):569–587, 2000.
[108] Joe H Ward Jr. Hierarchical grouping to optimize an objective function. *J. Amer. Statist. Assoc.*, 58(301):236–244, 1963.
[109] Neil White, G-C Rota, and Neil M White. *Theory of matroids*. Number 26. Cambridge University Press, 1986.
[110] V. Yurinsky. *Sums and Gaussian vectors*, volume 1617. Springer-Verlag, Berlin, 1995.
[111] Zhen-yue Zhang and Hong-yuan Zha. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *Journal of Shanghai University (English Edition)*, 8(4):406–424, 2004.
[112] GM Ziegler. Oriented matroids today: Dynamic survey and updated bibliography. *Electron. J. Combin*, 3, 1996.