



HAL
open science

Random sampling of bandlimited signals on graphs

Gilles Puy, Nicolas Tremblay, Rémi Gribonval, Pierre Vandergheynst

► **To cite this version:**

Gilles Puy, Nicolas Tremblay, Rémi Gribonval, Pierre Vandergheynst. Random sampling of bandlimited signals on graphs. 2015. hal-01229578v2

HAL Id: hal-01229578

<https://inria.hal.science/hal-01229578v2>

Preprint submitted on 20 May 2016 (v2), last revised 27 Jul 2016 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Random sampling of bandlimited signals on graphs

Gilles Puy, Nicolas Tremblay, Rémi Gribonval, and Pierre Vandergheynst

Abstract. We study the problem of sampling k -bandlimited signals on graphs. We propose two sampling strategies that consist in selecting a small subset of nodes at random. The first strategy is non-adaptive, *i.e.*, independent of the graph structure, and its performance depends on a parameter called the graph coherence. On the contrary, the second strategy is adaptive but yields optimal results. Indeed, no more than $O(k \log(k))$ measurements are sufficient to ensure an accurate and stable recovery of all k -bandlimited signals. This second strategy is based on a careful choice of the sampling distribution, which can be estimated quickly. Then, we propose a computationally efficient decoder to reconstruct k -bandlimited signals from their samples. We prove that it yields accurate reconstructions and that it is also stable to noise. Finally, we conduct several experiments to test these techniques.

1. Introduction

Graphs are a central modelling tool for network-structured data [1]. Depending on the application, the nodes of a graph may represent people in social networks, brain regions in neuronal networks, or stations in transportation networks. Data on a graph, such as individual hobbies, activity of brain regions, traffic at a station, may be represented by scalars defined on each node, which form a graph signal. Extending classical signal processing methods to graph signals is the purpose of the emerging field of graph signal processing [2, 3].

Within this framework, a cornerstone is sampling, *i.e.*, measuring a graph signal on a reduced set of nodes carefully chosen to enable stable reconstructions. Classically, sampling a continuous signal $x(t)$ consists in measuring a countable sequence of its values, $\{x(t_j)\}_{j \in \mathbb{Z}}$, that ensures its recovery under a given smoothness model [4]. Smoothness assumptions are often defined in terms of the signal's Fourier transform. For example, Shannon's famous sampling theorem [5] states that any ω -bandlimited signal can be recovered *exactly* from its values at $t_j = j/2\omega$. Similar theorems exist for other classes of signals, *e.g.*, signals on the sphere [6]; and other types of sampling schemes, *e.g.*, irregular sampling [7, 8] or compressive sampling [9]. Extending these theorems to graph signals requires to decide on a smoothness model and to design a sampling scheme that enables stable recovery.

Natural choices of smoothness models build upon, *e.g.*, the graph's adjacency matrix, the combinatorial Laplacian matrix, the normalised Laplacian, or the random walk Laplacian. The sets of eigenvectors of these operators define different graph Fourier bases. Given such a Fourier basis, the equivalent of a classical ω -bandlimited signal is a k -bandlimited graph signal whose k first Fourier coefficients are non-null [10, 11].

Unlike continuous time signal processing, the concept of regular sampling itself is not applicable for graph signals, apart for very regular graphs such as bipartite graphs [12]. We are left with two possible choices for sampling: irregular or random sampling. Irregular sampling of k -bandlimited graph signals has been studied first by Pesenson [13, 14] 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. Building upon this first work, and using the fact that the sampling matrix applied to the first k Fourier modes should have rank k in order to guarantee recovery of k -bandlimited signals, Anis *et al.* [11, 15] and Chen *et al.* [10, 16] showed that a sampling

G. Puy is with Technicolor, 975 Avenue des Champs Blancs, 35576 Cesson-Sévigné, France. N. Tremblay, R. Gribonval and P. Vandergheynst are with INRIA Rennes - Bretagne Atlantique, Campus de Beaulieu, FR-35042 Rennes Cedex, France. N. Tremblay and P. Vandergheynst are also with the Institute of Electrical Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland. The first version of this manuscript was submitted when G. Puy was at INRIA Rennes - Bretagne Atlantique, Campus de Beaulieu, FR-35042 Rennes Cedex, France. This work was partly funded by the European Research Council, PLEASE project (ERC-StG-2011-277906), and by the Swiss National Science Foundation, grant 200021-154350/1 - Towards Signal Processing on Graphs.

set of size k that perfectly embeds k -bandlimited signals always exists. To find such an optimal set, the authors need to compute the first k eigenvectors of the Laplacian, which is computationally prohibitive for large graphs. A recent work [17] bypasses the partial diagonalisation of the Laplacian by using graph spectral proxies, but the procedure to find an optimal sampling set still requires a search over all possible subsets of nodes of a given size. This is a very large combinatorial problem. In practice, approximate results are obtained using a greedy heuristic that enables the authors to efficiently perform experiments on graphs of size up to few thousands nodes.

Several other sampling schemes exist in the literature, such as schemes based on a bipartite decomposition of the graph [12], on a decomposition via maximum spanning trees [18], on the sign of the last Fourier mode [19], on the sign of the Fiedler vector [20], or on a decomposition in communities [21]. All these propositions are however specifically designed for graph multiresolution analysis with filterbanks, and are not suited to find optimal or close-to-optimal sets of nodes for sampling k -bandlimited graph signals.

1.1. Main Contributions

In this paper, we propose a very different approach to sampling on graphs. Instead of trying to find an *optimal* sampling set, (*i.e.*, a set of size k) for k -bandlimited signals, we relax this optimality constraint in order to tackle graphs of very large size. We allow ourselves to sample slightly more than k nodes and, inspired by compressive sampling, we propose two random sampling schemes that ensure recovery of graph signals with high probability.

A central graph characteristic that appears from our study is the *graph weighted coherence* of order k (see Definition 2.1). This quantity is a measure of the localisation of the first k Fourier modes on the nodes of the graph. Unlike the classical Fourier modes, some graph Fourier modes have the surprising potential of being localised on very few nodes. The farther a graph is from a regular grid, the higher the chance to have a few localised Fourier modes. This particularity in graph signal processing is studied in [22–24] but is still largely not understood.

First, we propose a non-adaptive sampling technique that consists in choosing a few nodes at random to form the sampling set. In this setting, we show that the number of samples ensuring the reconstruction of all k -bandlimited signals scales with the square of the graph weighted coherence. For regular or almost-regular graphs, *i.e.*, graphs whose coherence is close to \sqrt{k} , this result shows that $O(k \log k)$ samples selected using the uniform distribution are sufficient to sample k -bandlimited signals. We thus obtain an almost optimal sampling condition.

Second, for arbitrary graphs with a coherence potentially tending to \sqrt{n} , where $n \gg k$ is the total number of nodes, we propose a second sampling strategy that compensates the undesirable consequences of mode localisation. The technique relies on the variable density sampling strategy widely used in compressed sensing [25–27]. We prove that there always exists a sampling distribution such that no more than $O(k \log k)$ samples are sufficient to ensure exact and stable reconstruction of all k -bandlimited signals, whatever the graph structure. Unfortunately, computing the optimal sampling distribution requires the partial diagonalisation of the first k eigenvectors of the Laplacian. To circumvent this issue, we propose a fast technique to estimate this optimal sampling distribution accurately.

Finally, we propose an efficient method to reconstruct any k -bandlimited signal from its samples. We prove that the method recovers k -bandlimited signals exactly in the absence of noise. We also prove that the method is robust to measurement noise and model errors.

Note that our sampling theorems are applicable to any symmetrical Laplacian or adjacency matrix, *i.e.*, any weighted undirected graph. Nevertheless, the efficient recovery method we propose is specifically designed to take advantage of the semi-definite positivity of the Laplacian operator. In the following, we therefore concentrate on such symmetrical positive semi-definite Laplacians, such as the combinatorial or normalized Laplacians.

Let us acknowledge that the idea of random sampling for k -bandlimited graph signals is mentioned in [10] and [28]. In [10], the authors prove that the space of k -bandlimited graph signals can be stably embedded using a uniform sampling but for the Erdős-Rényi graph only. The idea of using a

non-uniform sampling appears in [28]. However, the authors do not prove that this sampling strategy provides a stable embedding of the space of k -bandlimited graph signals. We prove this result in Section 2 but also show that there always exists a sampling distribution that yields optimal results. Finally, the reconstruction methods proposed in [28] requires a partial diagonalisation of the Laplacian matrix, unlike ours. We also have much stronger recovery guarantees than the ones presented in [28], which are expected recovery guarantees.

1.2. Notations and definitions

For any matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, $\|\mathbf{X}\|_2$ denotes the spectral norm of \mathbf{X} , $\lambda_{\max}(\mathbf{X})$ denotes the largest eigenvalue of \mathbf{X} , and $\lambda_{\min}(\mathbf{X})$ denotes the smallest eigenvalue of \mathbf{X} . For any vector $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_2$ denotes the Euclidean norm of \mathbf{x} . Depending on the context, \mathbf{x}_j may represent the j^{th} entry of the vector \mathbf{x} or the j^{th} column-vector of the matrix \mathbf{X} . The identity matrix is denoted by \mathbf{I} - its dimensions are determined by the context - and $\boldsymbol{\delta}_j$ is its j^{th} column vector.

We consider an undirected, connected, weighted graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{W}\}$, where \mathcal{V} is the set of n nodes, \mathcal{E} is the set of edges, and $\mathbf{W} \in \mathbb{R}^{n \times n}$ is the weighted adjacency matrix. The entries of \mathbf{W} are nonnegative. We denote the graph Laplacian by $\mathbf{L} \in \mathbb{R}^{n \times n}$. As said before, we assume that \mathbf{L} is real, symmetric, and positive semi-definite. For example, the matrix \mathbf{L} can be the combinatorial graph Laplacian $\mathbf{L} := \mathbf{D} - \mathbf{W}$, or the normalised one $\mathbf{L} := \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$, where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is the diagonal degree matrix and \mathbf{I} is the identity matrix [29]. The diagonal degree matrix \mathbf{D} has entries $d_i := \sum_{i \neq j} \mathbf{W}_{ij}$.

As the matrix \mathbf{L} is real symmetric, there exists a set of orthonormal eigenvectors $\mathbf{U} \in \mathbb{R}^{n \times n}$ and real eigenvalues $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_n$ such that $\mathbf{L} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^\top$, where $\boldsymbol{\Lambda} := \text{diag}(\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_n) \in \mathbb{R}^{n \times n}$. Furthermore, semi-definite positivity of \mathbf{L} implies that all eigenvalues are nonnegative. Without loss of generality, we assume that $\boldsymbol{\lambda}_1 \leq \dots \leq \boldsymbol{\lambda}_n$.

The matrix \mathbf{U} is often viewed as the graph Fourier transform [2]. For any signal $\mathbf{x} \in \mathbb{R}^n$ defined on the nodes of the graph \mathcal{G} , $\hat{\mathbf{x}} = \mathbf{U}^\top \mathbf{x}$ contains the Fourier coefficients of \mathbf{x} ordered in increasing frequencies. As explained before, it is thus natural to consider that a k -bandlimited (smooth) signal $\mathbf{x} \in \mathbb{R}^n$ on \mathcal{G} with band-limit $k > 0$ is a signal that satisfies

$$\mathbf{x} = \mathbf{U}_k \hat{\mathbf{x}}^k$$

where $\hat{\mathbf{x}}^k \in \mathbb{R}^k$ and

$$\mathbf{U}_k := (\mathbf{u}_1, \dots, \mathbf{u}_k) \in \mathbb{R}^{n \times k},$$

i.e., \mathbf{U}_k is the restriction of \mathbf{U} to its first k vectors. This yields the following formal definition of a k -bandlimited signal.

Definition 1.1 (k -bandlimited signal on \mathcal{G}). *A signal $\mathbf{x} \in \mathbb{R}^n$ defined on the nodes of the graph \mathcal{G} is k -bandlimited with $k \in \mathbb{N} \setminus \{0\}$ if $\mathbf{x} \in \text{span}(\mathbf{U}_k)$.*

Note we use $\text{span}(\mathbf{U}_k)$ in our definition of k -bandlimited signals to handle the case where the eigendecomposition is not unique. To avoid any ambiguity in the definition of k -bandlimited signals, we assume that $\boldsymbol{\lambda}_k \neq \boldsymbol{\lambda}_{k+1}$ for simplicity.

1.3. Outline

In Section 2, we detail our sampling strategies and provide sufficient sampling conditions that ensure a stable embedding of k -bandlimited graph signals. We also prove that there always exists an optimal sampling distribution that ensures an embedding of k -bandlimited signals for $O(k \log(k))$ measurements. In Section 3, we propose decoders able to recover k -bandlimited signals from their samples. In Section 4, we explain how to obtain an estimation of the optimal sampling distribution quickly, *without* partial diagonalisation of the Laplacian matrix. In Section 5, we conduct several experiments on different graphs to test our methods. Finally, we conclude and discuss perspectives in Section 6.

2. Sampling k -bandlimited signals

In this section, we start by describing how we select a subset of the nodes to sample k -bandlimited signals. Then, we prove that this sampling procedure stably embeds the set of k -bandlimited signals. We describe how to reconstruct such signals from these measurements in Section 3.

2.1. The sampling procedure

In order to select the subset of nodes that will be used for sampling, we need a probability distribution \mathcal{P} on $\{1, \dots, n\}$. This probability distribution is used as a sampling distribution. We represent it by a vector $\mathbf{p} \in \mathbb{R}^n$. We assume that $\mathbf{p}_i > 0$ for all $i = 1, \dots, n$. We obviously have $\|\mathbf{p}\|_1 = \sum_{i=1}^n \mathbf{p}_i = 1$. We associate the matrix

$$\mathbf{P} := \text{diag}(\mathbf{p}) \in \mathbb{R}^{n \times n}$$

to \mathbf{p} .

The subset of nodes $\Omega := \{\omega_1, \dots, \omega_m\}$ used for sampling is constructed by drawing independently (with replacements) m indices from the set $\{1, \dots, n\}$ according to the probability distribution \mathbf{p} . We thus have

$$\mathbb{P}(\omega_j = i) = \mathbf{p}_i, \quad \forall j \in \{1, \dots, m\} \text{ and } \forall i \in \{1, \dots, n\}.$$

For any signal $\mathbf{x} \in \mathbb{R}^n$ defined on the nodes of the graph, its sampled version $\mathbf{y} \in \mathbb{R}^m$ satisfies

$$\mathbf{y}_j := \mathbf{x}_{\omega_j}, \quad \forall j \in \{1, \dots, m\}.$$

Note that we discuss the case of sampling without replacement in Section 2.3.

Let us pause for a moment and highlight few important facts. First, the sampling procedure allows each node to be selected multiple times. The number of measurements m includes these duplications. In practice, one can sample each selected node only once and add these duplications ‘‘artificially’’ afterwards. Second, *the set of nodes Ω needs to be selected only once to sample all k -bandlimited signals on \mathcal{G}* . One does not need to construct a set Ω each time a signal has to be sampled. Third, note that the sampling procedure is so far completely independent of the graph \mathcal{G} . This is a non-adaptive sampling strategy.

Let us define the sampling matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$. This matrix satisfies

$$(1) \quad \mathbf{M}_{ij} := \begin{cases} 1 & \text{if } j = \omega_i \\ 0 & \text{otherwise,} \end{cases}$$

for all $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n\}$. Note that $\mathbf{y} = \mathbf{M}\mathbf{x}$. In the next section, we show that, with high probability, \mathbf{M} embeds the set of k -bandlimited signals for a number of measurements m essentially proportional to $k \log(k)$ times a parameter called the graph weighted coherence.

2.2. The space of k -bandlimited signals is stably embedded

Similarly to many compressed sensing results, the number of measurements required to stably sample k -bandlimited signals will depend on a quantity, called the *graph weighted coherence*, that represents how the energy of these signals spreads over the nodes. Before providing the formal definition of this quantity, let us give an intuition of what it represents and why it is important.

Consider the signal $\delta_i \in \mathbb{R}^n$ with value 1 at node i and 0 everywhere else. This signal has its energy concentrated entirely at the i^{th} node. Compute $\mathbf{U}_k^T \delta_i$, *i.e.*, the first k Fourier coefficients of δ_i . The ratio

$$\frac{\|\mathbf{U}_k^T \delta_i\|_2}{\|\mathbf{U}^T \delta_i\|_2} = \frac{\|\mathbf{U}_k^T \delta_i\|_2}{\|\delta_i\|_2} = \|\mathbf{U}_k^T \delta_i\|_2$$

characterises how much the energy of δ_i is concentrated on the first k Fourier modes. This ratio varies between 0 and 1. When it is equal to 1, this indicates that there exists k -bandlimited signals whose energy is solely concentrated at the i^{th} node; not sampling the i^{th} node jeopardises the chance of reconstructing these signals. When this ratio is equal to 0, then no k -bandlimited signal has a part of its energy on the i^{th} node; one can safely remove this node from the sampling set. We thus see that

the quality of our sampling method will depend on the interplay between the sampling distribution \mathbf{p} and the quantities $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2$ for $i \in \{1, \dots, n\}$. Ideally, we should have \mathbf{p}_i large wherever $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2$ is large and \mathbf{p}_i small wherever $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2$ is small. The interplay between \mathbf{p}_i and $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2$ is characterised by the *graph weighted coherence*.

Definition 2.1 (Graph weighted coherence). *Let $\mathbf{p} \in \mathbb{R}^n$ represent a sampling distribution on $\{1, \dots, n\}$. The graph weighted coherence of order k for the pair $(\mathcal{G}, \mathbf{p})$ is*

$$\nu_{\mathbf{p}}^k := \max_{1 \leq i \leq n} \left\{ \mathbf{p}_i^{-1/2} \|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2 \right\}.$$

The quantity $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2$ is called the *local graph coherence at node i* .

Let us highlight two fundamental properties of $\nu_{\mathbf{p}}^k$. First, we have

$$\nu_{\mathbf{p}}^k \geq \sqrt{k}.$$

Indeed, as the columns of \mathbf{U}_k are normalised to 1, we have

$$k = \|\mathbf{U}_k\|_{\text{Frob}}^2 = \sum_{i=1}^n \|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2^2 = \sum_{i=1}^n \mathbf{p}_i \frac{\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2^2}{\mathbf{p}_i} \leq \max_{1 \leq i \leq n} \left\{ \frac{\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2^2}{\mathbf{p}_i} \right\} \cdot \sum_{i=1}^n \mathbf{p}_i = (\nu_{\mathbf{p}}^k)^2.$$

Second, $\nu_{\mathbf{p}}^k$ is a quantity that depends solely on \mathbf{p} and $\text{span}(\mathbf{U}_k)$. The choice of the basis for $\text{span}(\mathbf{U}_k)$ does not matter in the definition of $\nu_{\mathbf{p}}^k$. Indeed, it suffices to notice that $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2^2 = \|\mathbf{P}_k(\boldsymbol{\delta}_i)\|_2^2$, where $\mathbf{P}_k(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the orthogonal projection onto $\text{span}(\mathbf{U}_k)$, whose definition is independent of the choice of the basis of $\text{span}(\mathbf{U}_k)$. The graph weighted coherence is thus a characteristic of the interaction between the signal model, *i.e.*, $\text{span}(\mathbf{U}_k)$, and the sampling distribution \mathbf{p} .

We are now ready to introduce our main theorem which shows that $m^{-1}\mathbf{M}\mathbf{P}^{-1/2}$ satisfies a restricted isometry property on the space of k -bandlimited signals.

Theorem 2.2 (Restricted isometry property). *Let \mathbf{M} be a random subsampling matrix constructed as in (1) with the sampling distribution \mathbf{p} . For any $\delta, \epsilon \in (0, 1)$, with probability at least $1 - \epsilon$,*

$$(2) \quad (1 - \delta) \|\mathbf{x}\|_2^2 \leq \frac{1}{m} \left\| \mathbf{M}\mathbf{P}^{-1/2} \mathbf{x} \right\|_2^2 \leq (1 + \delta) \|\mathbf{x}\|_2^2$$

for all $\mathbf{x} \in \text{span}(\mathbf{U}_k)$ provided that

$$(3) \quad m \geq \frac{3}{\delta^2} (\nu_{\mathbf{p}}^k)^2 \log \left(\frac{2k}{\epsilon} \right).$$

Proof. See Appendix A. □

There are several important comments to make about the above theorem.

- First, this theorem shows that the matrix $\mathbf{M}\mathbf{P}^{-1/2}$ embeds the set of k -bandlimited signals into \mathbb{R}^m . Indeed, for any $\mathbf{x} \neq \mathbf{z} \in \text{span}(\mathbf{U}_k)$, we have $\left\| \mathbf{M}\mathbf{P}^{-1/2} (\mathbf{x} - \mathbf{z}) \right\|_2^2 \geq m(1 - \delta) \|\mathbf{x} - \mathbf{z}\|_2^2 > 0$, as $(\mathbf{x} - \mathbf{z}) \in \text{span}(\mathbf{U}_k)$. The matrix $\mathbf{M}\mathbf{P}^{-1/2}$ can thus be used to sample k -bandlimited signals.
- Second, we notice that $\mathbf{M}\mathbf{P}^{-1/2} \mathbf{x} = \mathbf{P}_\Omega^{-1/2} \mathbf{M}\mathbf{x}$ where $\mathbf{P}_\Omega \in \mathbb{R}^{m \times m}$ is the diagonal matrix with entries $(\mathbf{P}_\Omega)_{ii} = \mathbf{p}_{\omega_i}$. Therefore, one just needs to measure $\mathbf{M}\mathbf{x}$ in practice; the re-weighting by $\mathbf{P}_\Omega^{-1/2}$ can be done off-line.
- Third, as $(\nu_{\mathbf{p}}^k)^2 \geq k$, we need to sample at least k nodes. Note that k is also the minimum number of measurements that one must take to hope to reconstruct $\mathbf{x} \in \text{span}(\mathbf{U}_k)$.

The above theorem is quite similar to known compressed sensing results in bounded orthonormal systems [30, 31]. The proof actually relies on the same tools as the ones used in compressed sensing. However, in our case, the setting is simpler. Unlike in compressed sensing where the signal model is a union of subspaces, the model here is a single known subspace. In the proof, we exploit this fact to refine and tighten the sampling condition. In this simpler setting and thanks to our refined result, we can propose a sampling procedure that is *always* optimal in terms of the number of measurements.

In order to minimise the number of measurements, the idea is to choose a sampling distribution that minimises $\nu_{\mathbf{p}}^k$. Luckily, it occurs that it is always possible to reach the lower bound of $\nu_{\mathbf{p}}^k$ with a proper choice of the sampling distribution. The sampling distribution $\mathbf{p}^* \in \mathbb{R}^n$ that minimises the graph weighted coherence is

$$(4) \quad \mathbf{p}_i^* := \frac{\|\mathbf{U}_k^\top \delta_i\|_2^2}{k}, \quad i = 1, \dots, n,$$

for which $(\nu_{\mathbf{p}^*}^k)^2 = k$. The proof is simple. One just need to notice that $\sum_{i=1}^n \mathbf{p}_i^* = k^{-1} \sum_{i=1}^n \|\mathbf{U}_k^\top \delta_i\|_2^2 = k^{-1} \|\mathbf{U}_k\|_{\text{Frob}}^2 = k^{-1}k = 1$ so that \mathbf{p}^* is a valid probability distribution. Finally, it is easy to check that $(\nu_{\mathbf{p}^*}^k)^2 = k$. This yields the following corollary to Theorem 2.2.

Corollary 2.3. *Let \mathbf{M} be a random subsampling matrix constructed as in (1) with the sampling distribution \mathbf{p}^* defined in (4). For any $\delta, \epsilon \in (0, 1)$, with probability at least $1 - \epsilon$,*

$$(1 - \delta) \|\mathbf{x}\|_2^2 \leq \frac{1}{m} \left\| \mathbf{M} \mathbf{P}^{-1/2} \mathbf{x} \right\|_2^2 \leq (1 + \delta) \|\mathbf{x}\|_2^2$$

for all $\mathbf{x} \in \text{span}(\mathbf{U}_k)$ provided that

$$m \geq \frac{3}{\delta^2} k \log \left(\frac{2k}{\epsilon} \right).$$

The sampling distribution \mathbf{p}^* is optimal in the sense that the number of measurements needed to embed the set of k -bandlimited signals is essentially reduced to its minimum value. Note that, unlike Theorem 2.2 where the sampling is non-adaptive, the sampling distribution is now adapted to the structure of the graph and *a priori* requires the knowledge of a basis of $\text{span}(\mathbf{U}_k)$. We present a fast method that does not require the computation of a basis of $\text{span}(\mathbf{U}_k)$ to estimate \mathbf{p}^* in Section 4.

It is important to mention that variable density sampling techniques are also popular in compressed sensing to reduce the sampling rate. We have been inspired by the works in this field to develop our sampling technique on graphs. In compressed sensing, the high efficiency of variable density sampling was first observed empirically in magnetic resonance imaging where the goal was to speed up the acquisition by reducing the amount of acquired data [32]. Theoretical evidence of the efficiency of this technique then appeared in [25–27] as well as in [33–35] where additional measurement constraints and structured sparsity patterns are considered. There are similarities between the theoretical results existing in the compressed sensing literature and the ones presented in this work as we use similar proof techniques. However, we refine the proofs to take into account our specific signal model: bandlimited signals on graphs. One specificity of this setting is that there always exists a sampling distribution for which sampling $O(k \log(k))$ nodes is enough to capture all k -bandlimited signals. We recall that up to the log factor, one cannot hope to reduce the number of measurements much further. A second originality is that we can rapidly estimate this optimal distribution using fast filtering techniques on graphs (see Section 4).

Finally, we would like to highlight the similarity between the concept of local graph coherence and the concept of leverage scores of a matrix. Let \mathbf{G} be a matrix and \mathbf{V} be the left singular vectors of \mathbf{G} , the leverage scores related to the best rank- r approximation of \mathbf{G} are $l_i := \|\mathbf{V}_r^\top \delta_i\|_2^2$, where \mathbf{V}_r contains the r eigenvectors of \mathbf{G} with largest eigenvalues (see, *e.g.*, [36] for more details). The only difference between the leverage scores and the local graph coherences of \mathbf{L} is thus that the latter involve the smallest eigenvalues while the leverage scores involve the largest eigenvalues. For the normalised graph Laplacian $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$, one can notice that the leverage scores of $\mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$ correspond exactly to the local graph coherences of \mathbf{L} . In machine learning, the leverage scores have been used, *e.g.*, to improve the performance of randomised algorithms that solve overdetermined least-square problems [37–39] or that compute a low-rank approximation of a given matrix [38, 40–42]. These algorithms work by building a sketch of the matrix of interest from a subset of its rows and/or columns. The leverage scores represent the importance of each row/columns in the dataset. The optimised sampling distribution for the rows/columns is then constructed by normalising the vector of leverage scores, as we do it here with the local graph coherence. Note also that fast algorithms

to estimate the leverage scores have been developed in [36]. In the future, it would be interesting to compare this method with ours, which explicitly uses the graph structure in the computations.

2.3. Sampling without replacement

We have seen that the proposed sampling procedure allows one node to be sampled multiple times. In the case of a uniform sampling distribution, we can solve the issue by considering a sampling of the nodes without replacement and still prove that the set of k -bandlimited signals is stably embedded. We denote by $\boldsymbol{\pi} \in \mathbb{R}^n$ the uniform distribution on $\{1, \dots, n\}$, $\pi_i = 1/n$ for all $i \in \{1, \dots, n\}$.

Theorem 2.4. *Let \mathbf{M} be a random subsampling matrix constructed as in (1) with Ω built by drawing m indices $\{\omega_1, \dots, \omega_m\}$ from $\{1, \dots, n\}$ uniformly at random **without** replacement. For any $\delta, \epsilon \in (0, 1)$, with probability at least $1 - \epsilon$,*

$$(1 - \delta) \|\mathbf{x}\|_2^2 \leq \frac{n}{m} \|\mathbf{M}\mathbf{x}\|_2^2 \leq (1 + \delta) \|\mathbf{x}\|_2^2$$

for all $\mathbf{x} \in \text{span}(\mathbf{U}_k)$ provided that

$$m \geq \frac{3}{\delta^2} (\nu_{\boldsymbol{\pi}}^k)^2 \log\left(\frac{2k}{\epsilon}\right).$$

Proof. See Appendix A. □

The attentive reader will notice that, unfortunately, the condition on m is identical to the case where the sampling is done with replacement. This is because the theorem that we use to prove this result is obtained by “coming back” to sampling with replacement. Yet, we believe that it is still interesting to mention this result for applications where one wants to avoid any duplicated lines in the sampling matrix \mathbf{M} , which, for example, ensures that $\|\mathbf{M}\|_2 = 1$.

In the general case of non-uniform distributions, we are unfortunately not aware of any result allowing us to handle the case of a sampling without replacement. Yet it would be interesting to study this scenario more carefully in the future as sampling without replacement seems more natural for practical applications.

2.4. Intuitive links between a graph’s structure and its local coherence

Recall that the local graph coherence at node i reads $\|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2$. We give here some examples showing how this quantity changes for different graphs.

Consider first the d -dimensional grid with periodic boundary conditions. In this case, the eigenvectors of its Laplacian are simply the d -dimensional classical Fourier modes. For simplicity, we suppose that $\lambda_k \neq \lambda_{k+1}$. In this case, one can show that the local coherence is independent of i : $\forall i \in \mathcal{V} \quad \|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2 = \sqrt{k/n}$. The optimal probability \mathbf{p}^* is therefore uniform for the d -dimensional grid with periodic boundary conditions. Without the periodic boundary conditions, the optimal probability is mostly constant with an increase when going to the boundary nodes. An example in 1 dimension is shown in Fig. 4 for the path graph.

Let us now consider a graph made of k disconnected components of size n_1, \dots, n_k . We have $\sum_{j=1}^k n_j = n$. Considering the combinatorial graph Laplacian, one can show that a basis of $\text{span}(\mathbf{U}_k)$ is the concatenation of the indicator vectors of each component. Moreover, $\text{span}(\mathbf{U}_k)$ is the eigenspace associated to the eigenvalue 0. The local coherence associated to node i in component j is $1/\sqrt{n_j}$, and the probability to choose this node reads $p_i^* = 1/(kn_j)$. If all components have the same size, *i.e.*, $n_1 = \dots = n_k$, the optimal sampling is the uniform sampling. If the components have different sizes, the smaller is a component, the larger is the probability of sampling one of its nodes. With the optimal sampling distribution, each component is sampled with probability $1/k$, no matter the size of the component. The probability that each component is sampled at least once - a necessary condition for perfect recovery - is thus higher than when using uniform sampling. One may relax this strictly disconnected component example into a loosely defined community-structured graph, where k sets of nodes (forming a partition of the n nodes) are more connected with themselves than with

the rest of the graph. In this case, one also expects that the probability to sample a node is inversely proportional to the size of the community it belongs to.

3. Signal recovery

In the last section, we proved that it is possible to embed the space of k -bandlimited signals into \mathbb{R}^m using a sparse matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$. We now have to design a procedure to estimate accurately any $\mathbf{x} \in \text{span}(\mathbf{U}_k)$ from its, possibly *noisy*, m samples. Let us consider that the samples $\mathbf{y} \in \mathbb{R}^m$ satisfy

$$\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{n},$$

where $\mathbf{n} \in \mathbb{R}^m$ models a noise. Note that \mathbf{n} can be *any* vector \mathbb{R}^m . We do not restrict our study to a particular noise structure. The vector \mathbf{n} can be used to represent, *e.g.*, errors relative to the signal model or correlated noise. Such a recovery problem of a graph signal given few measured nodes and under a smoothness model is reminiscent to the litterature on semi-supervised learning. We link our approach to the appropriate literature in Section 3.3.

3.1. Standard decoder

In a situation where one knows a basis of $\text{span}(\mathbf{U}_k)$, the standard method to estimate \mathbf{x} from \mathbf{y} is to compute the best approximation to \mathbf{y} from $\text{span}(\mathbf{U}_k)$, *i.e.*, to solve

$$(5) \quad \min_{\mathbf{z} \in \text{span}(\mathbf{U}_k)} \left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M}\mathbf{z} - \mathbf{y}) \right\|_2.$$

Note that we introduced a weighting by the matrix $\mathbf{P}_\Omega^{-1/2}$ in (5) to account for the fact that $m^{-1} \mathbf{P}_\Omega^{-1/2} \mathbf{M} = m^{-1} \mathbf{M} \mathbf{P}_\Omega^{-1/2}$ satisfies the RIP, not \mathbf{M} alone. The following theorem proves that the solution of (5) is a faithful estimation of \mathbf{x} .

Theorem 3.1. *Let Ω be a set of m indices selected independently from $\{1, \dots, n\}$ using a sampling distribution $\mathbf{p} \in \mathbb{R}^n$, and \mathbf{M} be the sampling matrix associated to Ω (see (1)). Let $\epsilon, \delta \in (0, 1)$ and suppose that m satisfies (3). With probability at least $1 - \epsilon$, the following holds for all $\mathbf{x} \in \text{span}(\mathbf{U}_k)$ and all $\mathbf{n} \in \mathbb{R}^m$.*

i) Let \mathbf{x}^ be the solution of Problem (5) with $\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{n}$. Then,*

$$(6) \quad \|\mathbf{x}^* - \mathbf{x}\|_2 \leq \frac{2}{\sqrt{m(1-\delta)}} \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2.$$

ii) There exist particular vectors $\mathbf{n}_0 \in \mathbb{R}^m$ such that the solution \mathbf{x}^ of Problem (5) with $\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{n}_0$ satisfies*

$$(7) \quad \|\mathbf{x}^* - \mathbf{x}\|_2 \geq \frac{1}{\sqrt{m(1+\delta)}} \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n}_0 \right\|_2.$$

Proof. See Appendix B. □

We notice that in the absence of noise $\mathbf{x}^* = \mathbf{x}$, as desired. In the presence of noise, the upper bound on the error between \mathbf{x}^* and \mathbf{x} increases linearly with $\|\mathbf{P}_\Omega^{-1/2} \mathbf{n}\|_2$. For a uniform sampling, we have $\|\mathbf{P}_\Omega^{-1/2} \mathbf{n}\|_2 = \sqrt{n} \|\mathbf{n}\|_2$. For a non-uniform sampling, we may have $\|\mathbf{P}_\Omega^{-1/2} \mathbf{n}\|_2 \gg \sqrt{n} \|\mathbf{n}\|_2$ for some *particular* draws of Ω and noise vectors \mathbf{n} . Indeed, some weights \mathbf{p}_{ω_i} might be arbitrarily close to 0. Unfortunately, one cannot in general improve the upper bound in (6) as proved by the second part of the theorem with (7). Non-uniform sampling can thus be very sensitive to noise unlike uniform sampling. However, this is a worst case scenario. First, it is unlikely to draw an index ω_i where \mathbf{p}_{ω_i} is small by construction of the sampling procedure. Second,

$$\mathbb{E} \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2^2 = n \|\mathbf{n}\|_2^2,$$

so that $\|\mathbf{P}_\Omega^{-1/2}\mathbf{n}\|_2$ is not too large on average over the draw of Ω . Furthermore, in our numerical experiments, we noticed that we have $\min_i \mathbf{p}_i = 1/(\alpha^2 n)$, where $\alpha > 1$ is a small constant¹, for the optimal sampling distributions $\mathbf{p} = \mathbf{p}^*$ obtained in practice. This yields $\|\mathbf{P}_\Omega^{-1/2}\mathbf{n}\|_2 \leq \alpha \sqrt{n} \|\mathbf{n}\|_2$, which shows that non-uniform sampling is just slightly more sensitive to noise than uniform sampling in practical settings, with the advantage of reducing the number of measurements. Non-uniform sampling is thus still a beneficial solution.

We have seen a first method to estimate \mathbf{x} from its measurements. This method has however a major drawback: it requires the estimation of a basis of \mathbf{U}_k , which can be computationally very expensive for large graphs. To overcome this issue, we propose an alternative decoder which is computationally much more efficient. This algorithm uses techniques developed to filter graph signals rapidly. We thus briefly recall the principle of these filtering techniques.

3.2. Fast filtering on graphs

A filter is represented by a function $h : \mathbb{R} \rightarrow \mathbb{R}$ in the Fourier (spectral) domain. The signal \mathbf{x} filtered by h is

$$\mathbf{x}_h := \mathbf{U} \text{diag}(\hat{\mathbf{h}}) \mathbf{U}^\top \mathbf{x} \in \mathbb{R}^n,$$

where $\hat{\mathbf{h}} = (h(\lambda_1), \dots, h(\lambda_n))^\top \in \mathbb{R}^n$. Filtering thus consists in multiplying point by point the Fourier transform of \mathbf{x} with $\hat{\mathbf{h}}$ and then computing the inverse Fourier transform of the resulting signal.

According to the above definition, filtering *a priori* requires the knowledge of the matrix \mathbf{U} . To avoid the computation of \mathbf{U} , one can approximate the function h by a polynomial

$$\kappa(t) = \sum_{i=0}^d \alpha_i t^i$$

of degree d and compute \mathbf{x}_κ , which will approximate \mathbf{x}_h . This computation can be done rapidly as it only requires matrix-vector multiplications with \mathbf{L} , which is sparse in most applications. Indeed,

$$\mathbf{x}_\kappa = \mathbf{U} \text{diag}(\hat{\kappa}) \mathbf{U}^\top \mathbf{x} = \sum_{i=0}^d \alpha_i \mathbf{U} \text{diag}(\lambda_1^i, \dots, \lambda_n^i) \mathbf{U}^\top \mathbf{x} = \sum_{i=0}^d \alpha_i \mathbf{L}^i \mathbf{x}.$$

We let the reader refer to [43] for more information on this fast filtering technique.

To simplify notations, for any polynomial function $\kappa(t) = \sum_{i=0}^d \alpha_i t^i$ and any matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, we define

$$(8) \quad \kappa(\mathbf{A}) := \sum_{i=0}^d \alpha_i \mathbf{A}^i.$$

Remark that $\kappa(\mathbf{L}) = \mathbf{U} \kappa(\Lambda) \mathbf{U}^\top$.

3.3. Efficient decoder

Instead of solving (5), we propose to estimate \mathbf{x} by solving the following problem

$$(9) \quad \min_{\mathbf{z} \in \mathbb{R}^n} \left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M}\mathbf{z} - \mathbf{y}) \right\|_2^2 + \gamma \mathbf{z}^\top g(\mathbf{L})\mathbf{z},$$

where $\gamma > 0$ and $g : \mathbb{R} \rightarrow \mathbb{R}$ is a nonnegative and nondecreasing polynomial function. These assumptions on g imply that $g(\mathbf{L})$ is positive semi-definite - hence (9) is convex - and that $0 \leq g(\lambda_1) \leq \dots \leq g(\lambda_n)$.

The intuition behind decoder (9) is quite simple. Consider, for simplicity, that g is the identity. The regularisation term becomes $\mathbf{z}^\top \mathbf{L}\mathbf{z}$. Remember that a k -bandlimited signal is a signal that lives in the span of the first k eigenvector of \mathbf{U} , *i.e.*, where the eigenvalues of \mathbf{L} are the smallest. The regularisation term satisfies $\mathbf{z}^\top \mathbf{L}\mathbf{z} = (\mathbf{z}^\top \mathbf{U}) \Lambda (\mathbf{U}^\top \mathbf{z})$, where Λ is the diagonal matrix containing the eigenvalues of \mathbf{L} .

¹In the numerical experiments presented below, we have α smaller or equal to 3 in all cases tested with the optimal sampling distribution for the graphs presented in Fig. 1.

Therefore, this term penalises signals with energy concentrated at high frequencies more than signals with energy concentrated at low frequencies. In other words, this regularisation term favours the reconstruction of low-frequency signals, *i.e.*, signals approximately bandlimited. Notice also that one can recover the standard decoder defined in (5) by substituting the function $i_{\lambda_k} : \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$, defined as

$$i_{\lambda_k}(t) := \begin{cases} 0 & \text{if } t \in [0, \lambda_k], \\ +\infty & \text{otherwise,} \end{cases}$$

for g in (9).

We argue that solving (9) is computationally efficient because \mathbf{L} is sparse in most applications. Therefore, any method solving (9) that requires only matrix-vector multiplications with $g(\mathbf{L})$ can be implemented efficiently, as it requires multiplications with \mathbf{L} only (recall the definition of $g(\mathbf{L})$ in (8)). Examples of such methods are the conjugate gradient method or any gradient descend methods. Let us recall that one can find a solution to (9) by solving

$$(10) \quad (\mathbf{M}^\top \mathbf{P}_\Omega^{-1} \mathbf{M} + \gamma g(\mathbf{L})) \mathbf{z} = \mathbf{M}^\top \mathbf{P}_\Omega^{-1} \mathbf{y}.$$

The next theorem bounds the error between the original signal \mathbf{x} and the solution of (9).

Theorem 3.2. *Let Ω be a set of m indices selected independently from $\{1, \dots, n\}$ using a sampling distribution $\mathbf{p} \in \mathbb{R}^n$, \mathbf{M} be the sampling matrix associated to Ω (see (1)), and $M_{\max} > 0$ be a constant such that $\|\mathbf{M}\mathbf{P}^{-1/2}\|_2 \leq M_{\max}$. Let $\epsilon, \delta \in (0, 1)$ and suppose that m satisfies (3). With probability at least $1 - \epsilon$, the following holds for all $\mathbf{x} \in \text{span}(\mathbf{U}_k)$, all $\mathbf{n} \in \mathbb{R}^n$, all $\gamma > 0$, and all nonnegative and nondecreasing polynomial functions g such that $g(\lambda_{k+1}) > 0$.*

Let \mathbf{x}^* be the solution of (5) with $\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{n}$. Then,

$$(11) \quad \|\boldsymbol{\alpha}^* - \mathbf{x}\|_2 \leq \frac{1}{\sqrt{m(1-\delta)}} \left[\left(2 + \frac{M_{\max}}{\sqrt{\gamma g(\lambda_{k+1})}} \right) \|\mathbf{P}_\Omega^{-1/2} \mathbf{n}\|_2 + \left(M_{\max} \sqrt{\frac{g(\lambda_k)}{g(\lambda_{k+1})}} + \sqrt{\gamma g(\lambda_k)} \right) \|\mathbf{x}\|_2 \right],$$

and

$$(12) \quad \|\boldsymbol{\beta}^*\|_2 \leq \frac{1}{\sqrt{\gamma g(\lambda_{k+1})}} \|\mathbf{P}_\Omega^{-1/2} \mathbf{n}\|_2 + \sqrt{\frac{g(\lambda_k)}{g(\lambda_{k+1})}} \|\mathbf{x}\|_2,$$

where $\boldsymbol{\alpha}^* := \mathbf{U}_k \mathbf{U}_k^\top \mathbf{x}^*$ and $\boldsymbol{\beta}^* := (\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top) \mathbf{x}^*$.

Proof. See Appendix B. □

In the above theorem, $\boldsymbol{\alpha}^*$ is the orthogonal projection of \mathbf{x}^* onto $\text{span}(\mathbf{U}_k)$ and $\boldsymbol{\beta}^*$ onto the orthogonal complement of $\text{span}(\mathbf{U}_k)$. To obtain a bound on $\|\mathbf{x}^* - \mathbf{x}\|_2$, one can simply use the triangle inequality and the bounds (11) and (12).

In the absence of noise, we thus have

$$\|\mathbf{x}^* - \mathbf{x}\|_2 \leq \frac{1}{\sqrt{m(1-\delta)}} \left(M_{\max} \sqrt{\frac{g(\lambda_k)}{g(\lambda_{k+1})}} + \sqrt{\gamma g(\lambda_k)} \right) \|\mathbf{x}\|_2 + \sqrt{\frac{g(\lambda_k)}{g(\lambda_{k+1})}} \|\mathbf{x}\|_2.$$

If $g(\lambda_k) = 0$, we notice that we obtain a perfect reconstruction. Note that as g is supposed to be nondecreasing and nonnegative, $g(\lambda_k) = 0$ implies that we also have $g(\lambda_1) = \dots = g(\lambda_{k-1}) = 0$. If $g(\lambda_k) \neq 0$, the above bound shows that we should choose γ as close as possible to² 0 and seek to minimise the ratio $g(\lambda_k)/g(\lambda_{k+1})$ to minimise the upper bound on the reconstruction error. Notice that if $g(\mathbf{L}) = \mathbf{L}^l$, with $l \in \mathbb{N}^*$, then the ratio $g(\lambda_k)/g(\lambda_{k+1})$ decreases as l increases. Increasing the power of \mathbf{L} and taking γ sufficiently small to compensate the potential growth of $g(\lambda_k)$ is thus a simple solution to improve the reconstruction quality in the absence of noise.

²Notice that if \mathbf{y} is in the range of \mathbf{M} , then the solution of Problem (9) tends to the solution of $\min_{\mathbf{z}} \mathbf{z}^\top g(\mathbf{L}) \mathbf{z}$ s.t. $\mathbf{y} = \mathbf{M}\mathbf{z}$ in the limit where $\gamma \rightarrow 0^+$.

In the presence of noise, for a fixed function g , the upper bound on the reconstruction error is minimised for a value of γ proportional to $\|\mathbf{P}_\Omega^{-1/2} \mathbf{n}\|_2 / \|\mathbf{x}\|_2$. To optimise the result further, one should seek to have $g(\boldsymbol{\lambda}_k)$ as small as possible and $g(\boldsymbol{\lambda}_{k+1})$ as large as possible.

Decoder (9) has close links to several “decoders” used in the semi-supervised learning literature [44] that attempt to estimate the label of unlabeled data from a small number of labeled data, by supposing that the label functions are smooth either 1) in the data space or 2) in a suitable transformed space –using similarity kernels that define graphs modeling the underlying manifold for instance– [45, 46], or 3) in both [47, 48]. In our work, smoothness is defined solely using the graph (case 2) which we suppose given; there is no equivalent of a data space (case 1) on which to define another smoothness constraint. Nevertheless, other types of smoothness could be considered instead of the Laplacian smoothness $\mathbf{z}^\top g(\mathbf{L})\mathbf{z}$. For instance, one could decide to use an l_1 penalisation of the graph difference operator, as in [49], to allow the signal to depart from the smoothness prior at some nodes. The closest semi-supervised framework to our naive decoder (5) is found in [50], where the authors constrain the solution to be in $\text{span}(\mathbf{U}_k)$ without specifying precisely the value of k ; and the closest technique to our efficient decoder (9) is found in [46] even though their cost function has an additional term of the form $\sum_{i \notin \Omega} z_i^2$ compared to ours. Another decoding method may be found in [51], where the authors have a similar cost function to ours. However, they work in the data space (case 1 above) and try to optimise this cost function directly in this space, *i.e.*, without explicitly constructing and using a graph.

Even though we use similar decoders than in the semi-supervised learning literature, let us stress that an important difference is that we choose beforehand which nodes to sample/label. In this sense, our work also has connections with the literature in active learning [52], more precisely with the works that concentrate on the offline (all nodes to label are chosen from the start), single-batch (the nodes to sample are drawn simultaneously) selection problem [53], such as in [54–57]. Yet another connection may be found in the area of Gaussian Random Fields [58, 59]. The originality of our method compared to these works comes from our specific smoothness model ($\text{span}(\mathbf{U}_k)$) for which we devise an original sampling scenario which ensures stable reconstruction when coupled with the decoder (9) (see Theorem 3.2).

4. Estimation of the optimal sampling distribution

In this section, we explain how to estimate the optimal sampling distribution \mathbf{p}^* efficiently. This distribution is entirely defined by the values $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2^2$, $i = 1, \dots, n$ (see (4)). In order to be able to deal with large graphs and potentially large k , we want to avoid the computation of a basis of $\text{span}(\mathbf{U}_k)$ to estimate this distribution. Instead, we take another route that consists in filtering a small number of random signals. Note that the idea of filtering few random signals to estimate the number of eigenvalues of a Hermitian matrix in a given interval is already proposed and studied in [60]. We show here that this technique can be used to estimate \mathbf{p}^* .

For this estimation, we will need to use low-pass filters. For any $\lambda > 0$, the filter ideal low-pass filter $b_\lambda : \mathbb{R} \rightarrow \mathbb{R}$ with cut-off frequency λ satisfies

$$b_\lambda(t) = \begin{cases} 1 & \text{if } t \in [0, \lambda], \\ 0 & \text{otherwise.} \end{cases}$$

4.1. Principle of the estimation

We recall that our goal is to estimate $\|\mathbf{U}_k^\top \boldsymbol{\delta}_i\|_2^2$ for all $i \in \{1, \dots, n\}$. To understand how our method works, consider that $\boldsymbol{\lambda}_k$ is known for the moment. Let $\mathbf{r} \in \mathbb{R}^n$ be a vector with independent random entries that follow a standard normal distribution. By filtering \mathbf{r} with $b_{\boldsymbol{\lambda}_k}$, we obtain

$$\mathbf{r}_{b_{\boldsymbol{\lambda}_k}} = \mathbf{U} \text{diag}(\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_k, 0, \dots, 0) \mathbf{U}^\top \mathbf{r} = \mathbf{U}_k \mathbf{U}_k^\top \mathbf{r}.$$

The estimation of the optimal sampling distribution is based on the following property. The i^{th} entry of $\mathbf{r}_{b_{\lambda_k}}$ is

$$(\mathbf{r}_{b_{\lambda_k}})_i = \mathbf{r}_{b_{\lambda_k}}^T \boldsymbol{\delta}_i = \mathbf{r}^T \mathbf{U}_k \mathbf{U}_k^T \boldsymbol{\delta}_i,$$

and the mean of $(\mathbf{r}_{b_{\lambda_k}})_i^2$ satisfies

$$\mathbb{E}(\mathbf{r}_{b_{\lambda_k}})_i^2 = \boldsymbol{\delta}_i^T \mathbf{U}_k \mathbf{U}_k^T \mathbb{E}(\mathbf{r} \mathbf{r}^T) \mathbf{U}_k \mathbf{U}_k^T \boldsymbol{\delta}_i = \boldsymbol{\delta}_i^T \mathbf{U}_k \mathbf{U}_k^T \mathbf{U}_k \mathbf{U}_k^T \boldsymbol{\delta}_i = \boldsymbol{\delta}_i^T \mathbf{U}_k \mathbf{U}_k^T \boldsymbol{\delta}_i = \|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2.$$

This shows that $(\mathbf{r}_{b_{\lambda_k}})_i^2$ is an unbiased estimation of $\|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2$, the quantity we want to evaluate. Therefore, a possibility to estimate the optimal sampling distribution consists in filtering L random signals $\mathbf{r}^1, \dots, \mathbf{r}^L$ with the same distribution as \mathbf{r} and average $(\mathbf{r}_{b_{\lambda_k}}^1)_i^2, \dots, (\mathbf{r}_{b_{\lambda_k}}^L)_i^2$ for each $i \in \{1, \dots, n\}$. The next theorem shows that if λ_k is known, then $L \geq O(\log(n))$ random vectors are sufficient to have an accurate estimation of $\|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2$.

In the theorem below, we consider a realistic scenario where we filter the signals with a polynomial approximation of b_λ . This theorem shows how this approximation affects the estimation of $\|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2$. We denote the polynomial filter approximating b_λ by $c_\lambda : \mathbb{R} \rightarrow \mathbb{R}$. It satisfies

$$(13) \quad c_\lambda = b_\lambda + \hat{e}_\lambda,$$

where $\hat{e}_\lambda : \mathbb{R} \rightarrow \mathbb{R}$ models the approximation error. We define

$$\mathbf{E}_\lambda := \text{diag}(\hat{e}_\lambda(\boldsymbol{\lambda}_1), \dots, \hat{e}_\lambda(\boldsymbol{\lambda}_n)) \in \mathbb{R}^{n \times n}.$$

Theorem 4.1. *Let $\mathbf{r}^1, \dots, \mathbf{r}^L \in \mathbb{R}^n$ be L independent zero-mean Gaussian random vectors with covariance $L^{-1} \mathbf{I}$. Denote by $\mathbf{r}_{c_\lambda}^1, \dots, \mathbf{r}_{c_\lambda}^L \in \mathbb{R}^n$ the signals $\mathbf{r}^1, \dots, \mathbf{r}^L$ filtered by c_λ with $\lambda > 0$. Let j^* be the largest integer such that $\lambda_{j^*} \leq \lambda$. There exists an absolute constant $C > 0$ such for any $\epsilon, \delta \in (0, 1)$, with probability at least $1 - \epsilon$, the filtered signals satisfy*

$$(1 - \delta) \left| \|\mathbf{U}_{j^*}^T \boldsymbol{\delta}_i\|_2 - \|\mathbf{E}_\lambda \mathbf{U}^T \boldsymbol{\delta}_i\|_2 \right|^2 \leq \sum_{l=1}^L (\mathbf{r}_{c_\lambda}^l)_i^2 \leq (1 + \delta) \left| \|\mathbf{U}_{j^*}^T \boldsymbol{\delta}_i\|_2 + \|\mathbf{E}_\lambda \mathbf{U}^T \boldsymbol{\delta}_i\|_2 \right|^2,$$

for all $i \in \{1, \dots, n\}$, provided that

$$L \geq \frac{C}{\delta^2} \log \left(\frac{2n}{\epsilon} \right).$$

Proof. See Appendix C. □

The above theorem indicates that if $\lambda \in [\lambda_k, \lambda_{k+1})$ and \hat{e} is null, then $\sum_{l=1}^L (\mathbf{r}_{c_{\lambda_k}}^l)_i^2$ estimates $\|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2$ with an error at most δ on each entry $i \in \{1, \dots, n\}$. Recalling that the optimal sampling distribution has entries

$$\mathbf{p}_i^* = \frac{\|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2}{k} = \frac{\|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2}{\sum_{i=1}^n \|\mathbf{U}_k^T \boldsymbol{\delta}_i\|_2^2},$$

we see that $\tilde{\mathbf{p}} \in \mathbb{R}^n$ with entries

$$\tilde{\mathbf{p}}_i := \frac{\sum_{l=1}^L (\mathbf{r}_{c_{\lambda_k}}^l)_i^2}{\sum_{i=1}^n \sum_{l=1}^L (\mathbf{r}_{c_{\lambda_k}}^l)_i^2}$$

approximates the optimal sampling distribution. If we know λ_k and λ_{k+1} , we can thus approximate \mathbf{p}^* . In order to complete the method, we now need a solution to estimate λ_j with $j = k$ or $j = k + 1$.

4.2. Estimating λ_k and λ_{k+1}

Let $\lambda \in (0, \lambda_n)$. Theorem 4.1 shows that, with probability $1 - \epsilon$,

$$(1 - \delta) \sum_{i=1}^n \|\mathbf{U}_{j^*}^\top \boldsymbol{\delta}_i\|_2^2 \leq \sum_{i=1}^n \sum_{l=1}^L (\mathbf{r}_{b_\lambda}^l)_i^2 \leq (1 + \delta) \sum_{i=1}^n \|\mathbf{U}_{j^*}^\top \boldsymbol{\delta}_i\|_2^2,$$

when using the filter b_λ . Noticing that

$$\sum_{i=1}^n \|\mathbf{U}_{j^*}^\top \boldsymbol{\delta}_i\|_2^2 = \|\mathbf{U}_{j^*}\|_{\text{Frob}}^2 = j^*,$$

as the columns of \mathbf{U} are normalised, yields

$$(1 - \delta) j^* \leq \sum_{i=1}^n \sum_{l=1}^L (\mathbf{r}_{b_\lambda}^l)_i^2 \leq (1 + \delta) j^*.$$

In other words, the total energy of the filtered signals is tightly concentrated around j^* , which is the largest integer such that $\lambda_{j^*} \leq \lambda$. Therefore, the total energy of the filtered signals provides an estimation of the number of eigenvalues of \mathbf{L} that are below λ .

Using this phenomenon, one can obtain, by dichotomy, an interval $(\underline{\lambda}, \bar{\lambda})$ such that $k - 1$ eigenvalues are below $\underline{\lambda}$ and k eigenvalues are below $\bar{\lambda}$ and thus obtain an estimation of λ_k . The same procedure can be used to estimate λ_{k+1} . Note that we cannot filter the signals using an ideal low-pass filter in practice, so that an additional error will slightly perturb the estimation.

4.3. The complete algorithm

We now have all the tools to design an algorithm that estimates the optimal sampling distribution. This is summarised in Algorithm 1. In practice, we noticed that using $L = 2 \log(n)$ signals is usually enough to obtain a reasonable approximation of the sampling distribution. We also only estimate λ_k and do not estimate λ_{k+1} . Steps 3 to 10 of the algorithm concern the estimation of λ_k by dichotomy. The estimated optimal sampling distribution $\tilde{\mathbf{p}} \in \mathbb{R}^n$ is defined in Step 10. Finally, we would like to mention that a better estimation of λ_k and of the sampling distribution could be obtained by running multiple times Algorithm 1 and averaging the results. In the following experiments, this algorithm is run only once but already yields good results.

Algorithm 1 Estimation of the optimal sampling distribution

Input: Precision parameter $\epsilon \in (0, 1)$, and bandlimit k .

- 1: Set $L = 2 \log(n)$ and draw L random vectors $\mathbf{r}^1, \dots, \mathbf{r}^L \in \mathbb{R}^n$ as in Theorem 4.1.
- 2: Estimate λ_n and set $\underline{\lambda} = 0$, $\bar{\lambda} = \lambda_n$, $\lambda = \lambda_n/2$, and compute c_λ that approximates the ideal low-pass filter b_λ .
- 3: **while** $\text{round}\left(\sum_{i=1}^n \sum_{l=1}^L (\mathbf{r}_{c_\lambda}^l)_i^2\right) \neq k$ or $|\underline{\lambda} - \bar{\lambda}| > \epsilon \cdot \bar{\lambda}$ **do**
- 4: **if** $\text{round}\left(\sum_{i=1}^n \sum_{l=1}^L (\mathbf{r}_{c_\lambda}^l)_i^2\right) \geq k$ **then**
- 5: Set $\bar{\lambda} = \lambda$.
- 6: **else**
- 7: Set $\underline{\lambda} = \lambda$.
- 8: **end if**
- 9: Set $\lambda = (\underline{\lambda} + \bar{\lambda})/2$, and compute c_λ that approximates the ideal low-pass filter b_λ .
- 10: **end while**

Output: Set $\tilde{\mathbf{p}}_i = \left(\sum_{l=1}^L (\mathbf{r}_{c_\lambda}^l)_i^2\right) / \left(\sum_{i=1}^n \sum_{l=1}^L (\mathbf{r}_{c_\lambda}^l)_i^2\right)$.

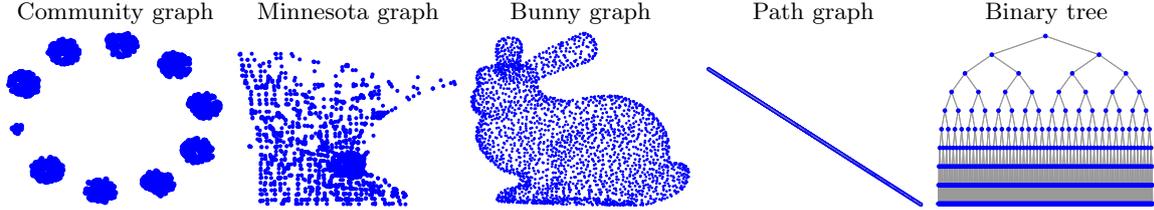


Figure 1. The five different graphs used in the simulations.

5. Experiments

In this section, we run several experiments to illustrate the above theoretical findings. First we show how the sampling distribution affects the number of measurements required to ensure that the RIP holds. Then, we show how the reconstruction quality is affected with the choice of g and γ in (9).

All our experiments are done using five different types of graph, all available in the GSP toolbox [61] and presented in Fig. 1. We use a) different community-type graphs of size $n = 1000$, b) the graph representing the Minnesota road network of size $n = 2642$, c) the graph of the Stanford bunny of size $n = 2503$, d) the unweighted path graph of size $n = 1000$ and e) a binary tree of depth 9 and size $n = 1023$. We recall that each node in the unweighted path graph is connected to its left and right neighbours with weight 1, except for the two nodes at the boundaries which have only one neighbour. We use the combinatorial Laplacian in all experiments. All samplings are done in the conditions of Theorem 2.2, *i.e.*, with replacement. Finally, the reconstructions are obtained by solving (10) using the `mldivide` function of Matlab. For the graphs and functions g considered, we noticed that it was faster to use this function than solving (10) by conjugate gradient.

5.1. Effect of the sampling distribution on m

In this first part, we study how the sampling distributions affect the minimum number of measurements required to satisfy the RIP. All experiments are repeated for three different sampling distributions: a) the uniform distribution $\boldsymbol{\pi}$, b) the optimal distribution \boldsymbol{p}^* , and c) the estimated optimal distribution $\tilde{\boldsymbol{p}} \in \mathbb{R}^n$ computed using Algorithm 1.

5.1.1. Using community graphs

We conduct a first set of experiments using five types of community graph, denoted by $\mathcal{C}_1, \dots, \mathcal{C}_5$. They all have 10 communities. To study the effect of the size of the communities on the sampling distribution, we choose to build these graphs with 9 communities of (approximately) equal size and reduce the size of last community:

- the graphs of type \mathcal{C}_1 have 10 communities of size 100;
- the graphs of type \mathcal{C}_2 have 1 community of size 50, 8 communities of size 105, and 1 community of size 110;
- the graphs of type \mathcal{C}_3 have 1 community of size 25, 8 communities of size 108, and 1 community of size 111;
- the graphs of type \mathcal{C}_4 have 1 community of size 17, 8 communities of size 109, and 1 community of size 111;
- the graphs of type \mathcal{C}_5 have 1 community of size 13, 8 communities of size 109, and 1 community of size 115.

For each pair of graph-type, $j \in \{1, \dots, 5\}$, and sampling distribution, $\boldsymbol{p} \in \{\boldsymbol{\pi}, \boldsymbol{p}^*, \tilde{\boldsymbol{p}}\}$, we generate a graph of type \mathcal{C}_j , compute \mathbf{U}_{10} and the lower RIP constant

$$(14) \quad \underline{\delta}_{10} := 1 - \inf_{\substack{\boldsymbol{x} \in \text{span}(\mathbf{U}_{10}) \\ \|\boldsymbol{x}\|_2 = 1}} \left\{ \frac{1}{m} \left\| \text{MP}^{-1/2} \boldsymbol{x} \right\|_2^2 \right\},$$

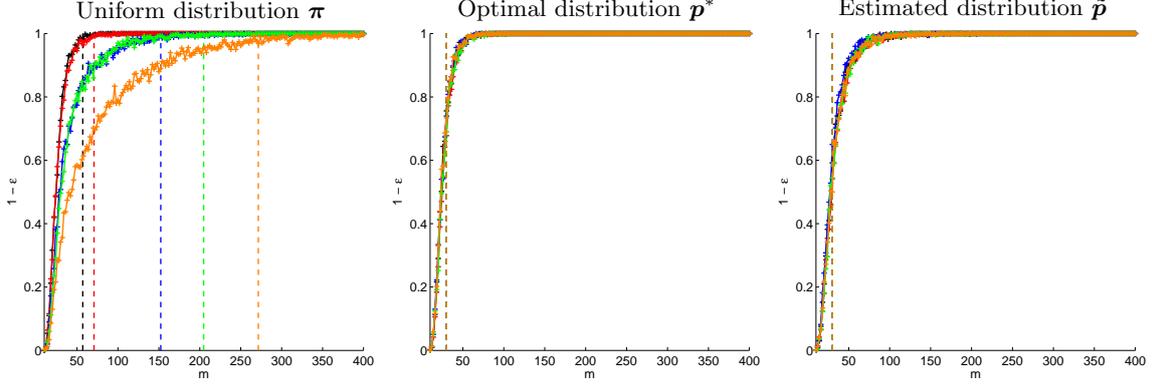


Figure 2. Probability that $\underline{\delta}_{10}$ is less than 0.995 as a function of m for 5 different types of community graph: C_1 in black, C_2 in red, C_3 in blue, C_4 in green, C_5 in orange. Left panel: the dashed vertical lines indicate the value of $3 \cdot (\nu_{\pi}^{10})^2$ for each type of graph. Middle and right panels: the dashed vertical lines indicate the value $3 \cdot (\nu_{\mathbf{p}^*}^{10})^2 = 3 \cdot 10$.

for different numbers of measurements m . Note that to compute $\underline{\delta}_{10}$, one just needs to notice that

$$\underline{\delta}_{10} = 1 - \frac{1}{m} \lambda_{\min} \left(\mathbf{U}_{10}^T \mathbf{P}^{-1/2} \mathbf{M}^T \mathbf{M} \mathbf{P}^{-1/2} \mathbf{U}_{10} \right).$$

We compute $\underline{\delta}_{10}$ for 500 independent draws of the matrix \mathbf{M} . When conducting the experiments with the estimated optimal distribution $\tilde{\mathbf{p}}$, we re-estimate this distribution at each of the 500 trials.

We present in Fig. 2 the probability that $\underline{\delta}_{10}$ is less than 0.995, estimated over the 500 trials, as a function of m . Let $m_{j,\mathbf{p}}^*$ be the number of measurements required to reach a probability of, e.g., $\mathbb{P}(\underline{\delta}_{10} \leq 0.995) = 0.9$ for the pair (j, \mathbf{p}) of graph-type and sampling distribution. Theorem 2.2 predicts that $m_{j,\mathbf{p}}^*$ scales linearly with $(\nu_{\mathbf{p}}^{10})^2$.

- For the uniform distribution π , the first figure from the left in Fig. 2 indicates the value of $(\nu_{\pi}^{10})^2(j)$, $j = 1, \dots, 5$ for the five different types of graph. We have $(\nu_{\pi}^{10})^2(1) \leq \dots \leq (\nu_{\pi}^{10})^2(5)$ and $m_{1,\pi}^* \leq m_{2,\pi}^* \leq \dots \leq m_{5,\pi}^*$, in accordance with Theorem 2.2.
- For the optimal sampling distribution \mathbf{p}^* , we have $(\nu_{\mathbf{p}^*}^{10})^2 = 10$. Therefore m_{j,\mathbf{p}^*}^* must be identical for all graph-types, as observed in the second panel of Fig. 2.
- For the estimated optimal sampling distribution $\tilde{\mathbf{p}}$, the last figure in Fig. 2 shows that the performance is identical for all graph-types, as with \mathbf{p}^* . Furthermore, we attained almost the same performance with $\tilde{\mathbf{p}}$ and \mathbf{p}^* , confirming the quality of the estimation provided by Algorithm 1

5.1.2. Using the Minnesota and bunny graphs

To confirm the results observed above, we repeat the same experiments but using four other graphs: the Minnesota, the bunny and path graphs, and the binary tree. For the first three graphs, the experiments are performed for k -bandlimited signals with band-limits 10 and 100, *i.e.*, we compute $\underline{\delta}_k$ - defined as in (14) - with \mathbf{U}_k . For the binary tree, we set the band-limits at 16 and 64. These choices are due to the fact that some eigenvalues have a multiplicity larger than 1 for this last graph. These choices ensure that $\lambda_k < \lambda_{k+1}$, as required in our assumptions.

We present in Fig. 3 the probability that $\underline{\delta}_k$ is less than 0.995, estimated over 500 draws of \mathbf{M} , as a function of m .

For the Minnesota, the bunny and path graphs at $k = 10$, we remark that all distributions yield essentially the same result. The advantage of using the distributions \mathbf{p}^* or $\tilde{\mathbf{p}}$ is more obvious at $k = 100$ for the Minnesota and the bunny graphs. Note that for the bunny graph, we reach only a probability of 0.036 at $m = 2000$ with the uniform distribution, whereas $m = 600$ measurements are sufficient to reach a probability 1 with \mathbf{p}^* . Uniform sampling is not working for the bunny graph at

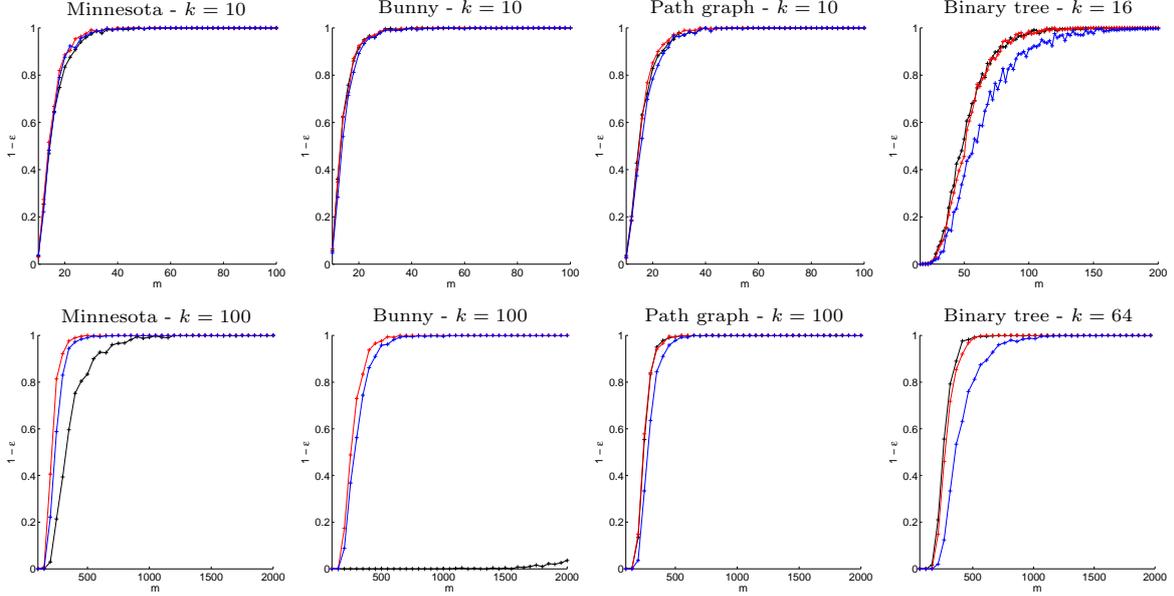


Figure 3. Probability that δ_k is less than 0.995 as a function of m . The curve in black indicates the result for the uniform distribution. The curve in red indicates the result for the optimal distribution. The curve in blue indicates the result for the estimated optimal distribution. The panels on the top row show the results at $k = 10$ for the Minnesota, bunny and path graphs, and at $k = 16$ for the binary tree. The panels on the bottom row show the results at $k = 100$ for the Minnesota, bunny and path graphs, and at $k = 64$ for the binary tree.

$k = 100$ because there exist few eigenmodes whose energy is highly concentrated on few nodes. In other words, we have $\|\mathbf{U}_{100}^T \delta_i\|_2 \approx 1$ for few nodes i . Finally, for the path graph at $k = 100$ and for the binary tree, we notice that the uniform distribution and the optimal distribution \mathbf{p}^* have the same performance while the estimated optimal distribution $\tilde{\mathbf{p}}$ performs slightly worse. However, we noticed that these differences decrease when increasing the number L of random signals in Algorithm (1).

5.1.3. Examples of optimal and estimated sampling distributions

For illustration, we present some examples of sampling distributions in Fig. 4 for five of the graphs used above. The top panels in Fig. 4 show the optimal sampling distribution computed with \mathbf{U}_k . The bottom panels show the estimated sampling distribution obtained with Algorithm 1.

For the community, Minnesota and bunny graphs, we notice that the estimated sampling distribution $\tilde{\mathbf{p}}$ and the optimal one \mathbf{p}^* are quite similar. We observe more differences for the path graph and the binary tree. This explains the slight differences of performance in the previous experiments. We recall that these differences decrease when increasing the number L of random signals in Algorithm 1.

It is interesting to notice that for the path graph, the optimal sampling distribution is essentially constant except for the nodes at the boundaries that are less connected than the other nodes and need to be sampled with higher probability. For the binary tree, the probability of sampling a node is only determined by its depth in the tree - all nodes at a given depth are equally important - and the deeper the node is, the higher the probability of sampling this node should be - it is easier to predict the value at one node from the values bore by its children than its parents.

5.2. Reconstruction of k -bandlimited signals

In this second part, we study experimentally the performance of the decoder (9). All experiments are repeated for 3 different graphs: a community graph of type \mathcal{C}_5 , the Minnesota graph and the bunny graph. We consider the recovery of k -bandlimited signals with band-limit $k = 10$. We take $m = 200$ measurements using the estimated optimal distribution $\tilde{\mathbf{p}}$. The experiments are conducted

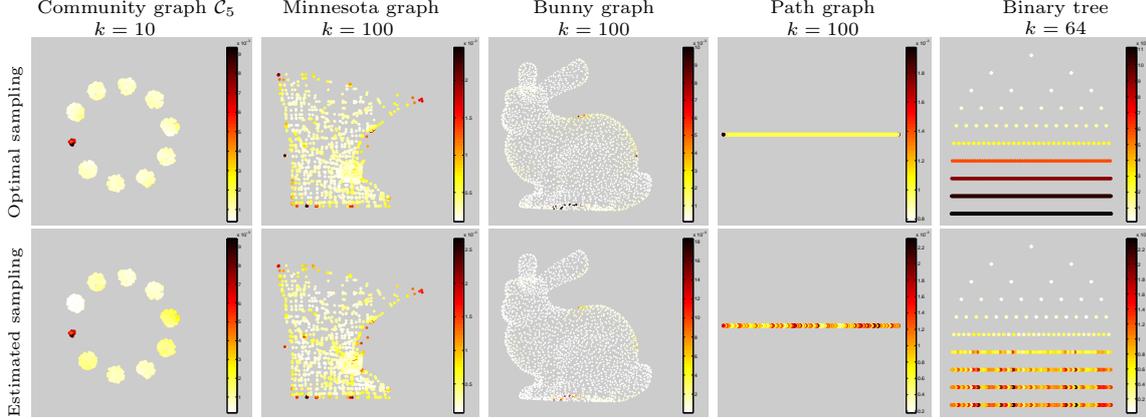


Figure 4. *Optimal and estimated optimal sampling distributions for five different graphs. Nodes in black are sampled with a higher probability than nodes in white.*

with and without noise on the measurements. In the presence of noise, the random noise vector \mathbf{n} follows a zero-mean Gaussian distribution³ of variance σ^2 . The values of σ used are $\{0, 1.5 \cdot 10^{-3}, 3.7 \cdot 10^{-3}, 8.8 \cdot 10^{-3}, 2.1 \cdot 10^{-2}, 5.0 \cdot 10^{-2}\}$. The signals are reconstructed by solving (9) for different values of the regularisation parameter γ and different functions g . For the community graph and the bunny graph, the regularisation parameter γ varies between 10^{-3} and 10^2 . For the Minnesota graph, it varies between 10^{-1} and 10^{10} . For each σ , 10 independent random signals of unit norm are drawn, sampled and reconstructed using all possible pairs (γ, g) . Then, we compute the mean reconstruction errors⁴ $\|\mathbf{x}^* - \mathbf{x}\|_2$, $\|\boldsymbol{\alpha}^* - \boldsymbol{\alpha}\|_2$ and $\|\boldsymbol{\beta}^*\|_2$ over these 10 signals. In our experiments, the distribution $\tilde{\mathbf{p}}$ is re-estimated with Algorithm 1 each time a new signal \mathbf{x} is drawn.

We present the mean reconstruction errors obtained in the absence of noise on the measurements in Fig. 5. In this set of experiments, we reconstruct the signals using $g(\mathbf{L}) = \mathbf{L}$, then $g(\mathbf{L}) = \mathbf{L}^2$, and finally $g(\mathbf{L}) = \mathbf{L}^4$. Before describing these results, we recall that the ratio $g(\lambda_{10})/g(\lambda_{11})$ decreases as the power of \mathbf{L} increases. We observe that all reconstruction errors, $\|\mathbf{x}^* - \mathbf{x}\|_2$, $\|\boldsymbol{\alpha}^* - \boldsymbol{\alpha}\|_2$ and $\|\boldsymbol{\beta}^*\|_2$ decrease when the ratio $g(\lambda_k)/g(\lambda_{k+1})$ in the range of small γ , as predicted by the upper bounds on these errors in Theorem 3.2.

We present the mean reconstruction errors obtained in the presence of noise on the measurements in Fig. 6. In this set of experiments, we reconstruct the signals using $g(\mathbf{L}) = \mathbf{L}^4$. As expected the best regularisation parameter γ increases with the noise level.

5.3. Illustration: sampling of a real image

We finish this experimental section with an example of image sampling using the developed theory.

For this illustration, we use the photo of *Lac d'Emosson* in Switzerland presented in Fig. 7(a) This RGB image contains 4288×2848 pixels. We divide this image into patches of 8×8 pixels, thus obtaining 536×356 patches of 64 pixels per RGB channel. Let us denote each patch by $\mathbf{q}_{i,j,l} \in \mathbb{R}^{64}$ with $i \in \{1, \dots, 536\}$, $j \in \{1, \dots, 356\}$, and $l \in \{1, 2, 3\}$. The pair of indices (i, j) encodes the spatial location of the patch and l encodes the color channel. Using these patches, we build the following matrix

$$\mathbf{X} := \begin{pmatrix} \mathbf{q}_{1,1,1} & \mathbf{q}_{1,2,1} & \cdots & \mathbf{q}_{2,1,1} & \cdots & \mathbf{q}_{536,356,1} \\ \mathbf{q}_{1,1,2} & \mathbf{q}_{1,2,2} & \cdots & \mathbf{q}_{2,1,2} & \cdots & \mathbf{q}_{536,356,2} \\ \mathbf{q}_{1,1,3} & \mathbf{q}_{1,2,3} & \cdots & \mathbf{q}_{2,1,3} & \cdots & \mathbf{q}_{536,356,3} \end{pmatrix} \in \mathbb{R}^{192 \times n},$$

where $n = 190816$. Each column of \mathbf{X} represents a color patch of the original image at a given position.

³For nodes sampled multiple times, the realisation of the noise is thus different each time the same node is sampled. The noise vector \mathbf{n} contains no duplicated entry.

⁴See Theorem 3.2 for the definition of $\boldsymbol{\alpha}^*$ and $\boldsymbol{\beta}^*$.

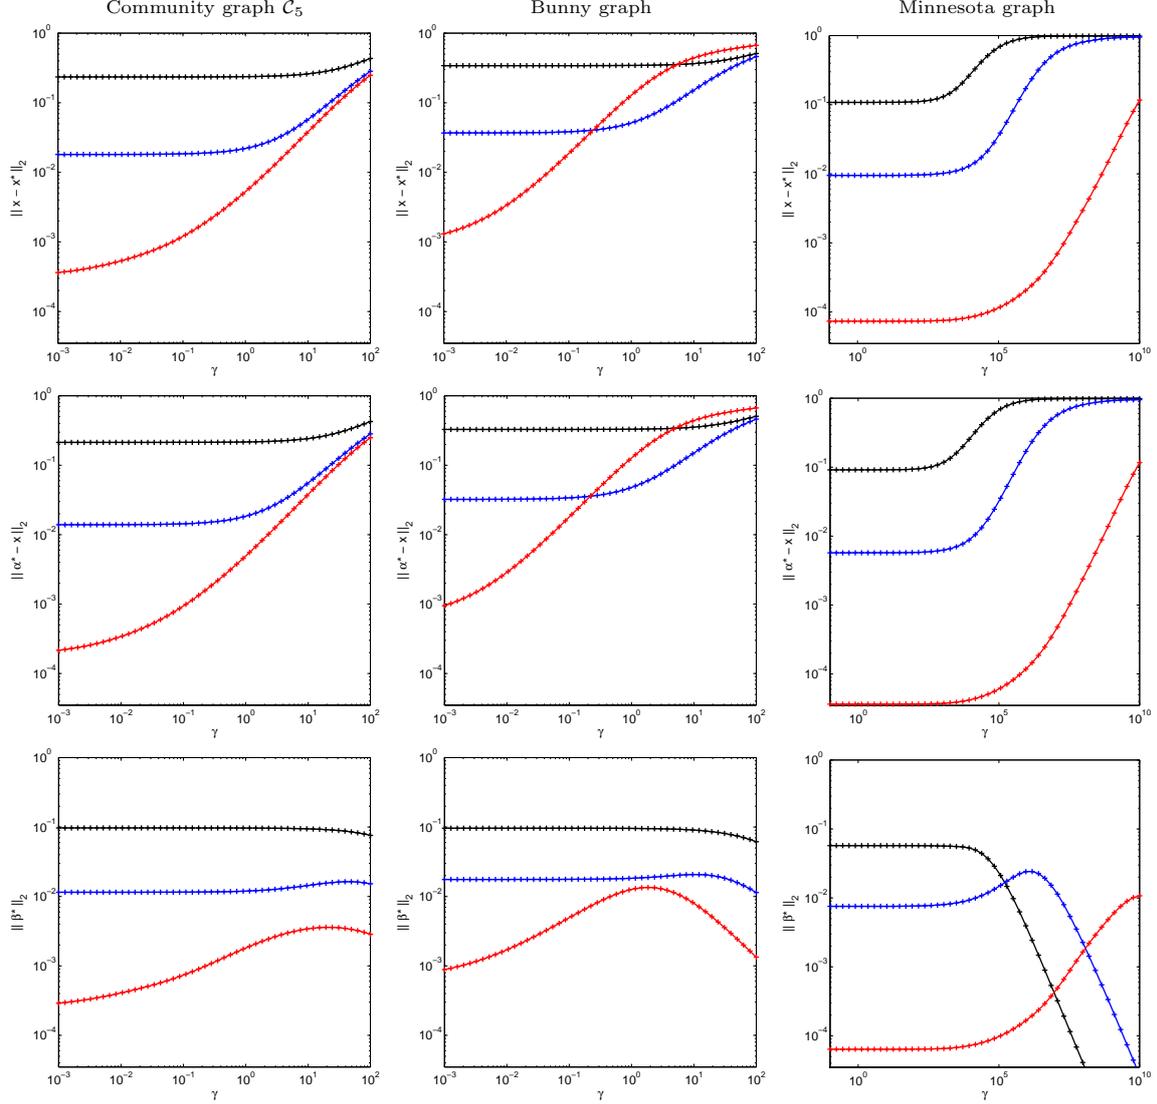


Figure 5. Mean reconstruction errors of 10-bandlimited signals as a function of γ . The simulations are performed in the absence of noise. The black curves indicate the results with $g(\mathbf{L}) = \mathbf{L}$. The blue curves indicate the results with $g(\mathbf{L}) = \mathbf{L}^2$. The red curves indicate the results with $g(\mathbf{L}) = \mathbf{L}^4$. The first, second and third columns show the results for a community graph of type \mathcal{C}_5 , the bunny graph, and the Minnesota graph, respectively. The first, second and third rows show the mean reconstruction errors $\|\mathbf{x}^* - \mathbf{x}\|_2$, $\|\boldsymbol{\alpha}^* - \mathbf{x}\|_2$ and $\|\boldsymbol{\beta}^*\|_2$, respectively.

We continue by building a graph modelling the similarity between the columns of \mathbf{X} . Let $\mathbf{x}_i \in \mathbb{R}^{192}$ be the i^{th} column-vector of the matrix \mathbf{X} . For each $i \in \{1, \dots, n\}$, we search for the 20 nearest neighbours of \mathbf{x}_i among all other columns of \mathbf{X} . Let $\mathbf{x}_j \in \mathbb{R}^{192}$ be a vector connected to \mathbf{x}_i . The weight W_{ij} of the weighted adjacency matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ satisfies

$$W_{ij} := \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\sigma^2}\right),$$

where $\sigma > 0$ is the standard deviation of all Euclidean distances between pairs of connected columns (patches). We then symmetrise the matrix \mathbf{W} . Each column of \mathbf{X} is thus connected to at least 20 other

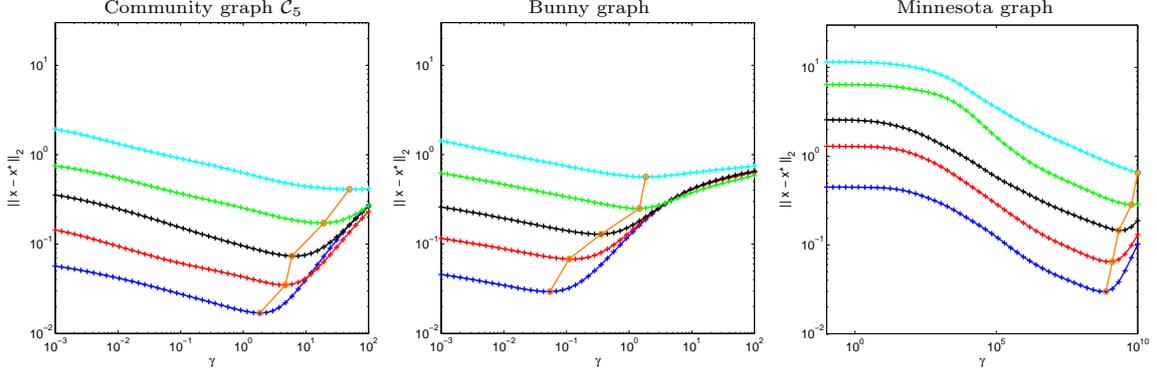


Figure 6. Mean reconstruction error $\|\mathbf{x}^* - \mathbf{x}\|_2$ of 10-bandlimited signals as a function of γ with $g(\mathbf{L}) = \mathbf{L}^4$. The simulations are performed in presence of noise. The standard deviation of the noise is 0.0015 (blue), 0.0037 (red), 0.0088 (black), 0.0210 (green), 0.0500 (cyan). The best reconstruction errors are indicated by orange circles. The first, second and third columns show the results for a community graph of type \mathcal{C}_5 , the bunny graph, and the Minnesota graph, respectively.

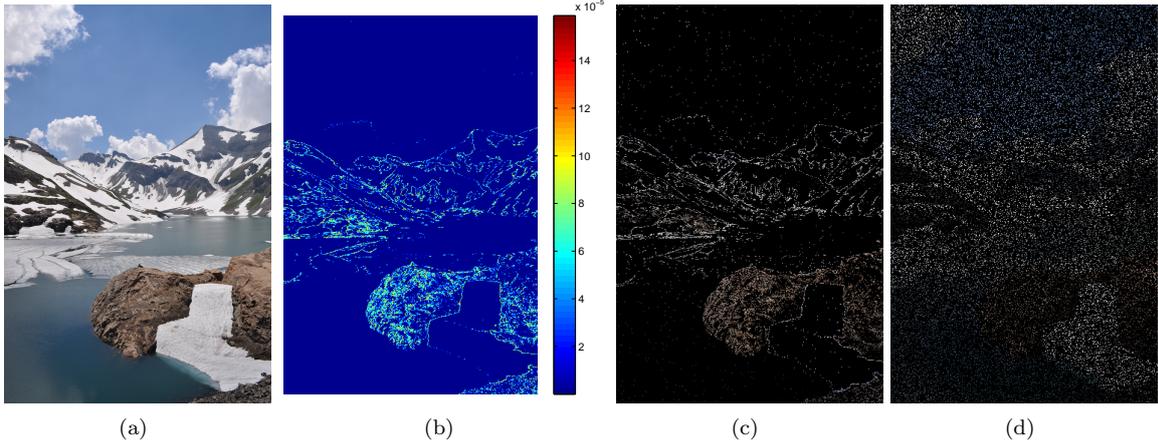


Figure 7. a) original image; b) estimated optimal sampling distribution $\tilde{\mathbf{p}}$; c) sampled image using $\tilde{\mathbf{p}}$; d) sampled image using the uniform sampling distribution $\boldsymbol{\pi}$. The sampled images are obtained using the same number of measurements.

columns after symmetrisation. We finish the construction of the graph by computing the combinatorial Laplacian $\mathbf{L} \in \mathbb{R}^{n \times n}$ associated to \mathbf{W}_{ij} .

We sample \mathbf{X} by measuring about 15% of its n columns: $m = 28622 \approx 0.15n$. First, we estimate the optimal sampling distribution $\tilde{\mathbf{p}}$ for $k = 9541 \approx m/3$ with Algorithm 1. It takes about 4 minutes to compute $\tilde{\mathbf{p}}$ using Matlab on a laptop with a 2,8 GHz Intel Core i7 and 16GB of RAM. In comparison, we tried to compute \mathbf{p}^* exactly by computing \mathbf{U}_{9541} but stopped Matlab after 30 minutes of computations. The estimated sampling distribution is presented in Fig. 7(b). Then, we build the sampling matrix \mathbf{M} by drawing at random m independent indices from $\{1, \dots, n\}$ according to $\tilde{\mathbf{p}}$. Note that the *effective sampling rate*, *i.e.*, once the indices sampled multiple times are removed, is about 7.6%, only. The sampled columns are denoted by $\mathbf{Y} \in \mathbb{R}^{192 \times m}$ and satisfy $\mathbf{Y}\boldsymbol{\tau} = \mathbf{M}\mathbf{X}\boldsymbol{\tau}$.

We present in Fig. 7(c) the sampled image, where all non-sampled pixels appear in black. We remark that the regions where many patches are similar (sky, lake, snow) are very sparsely sampled. This can be explained as follows. The patches in such a region being all similar, one can fill this

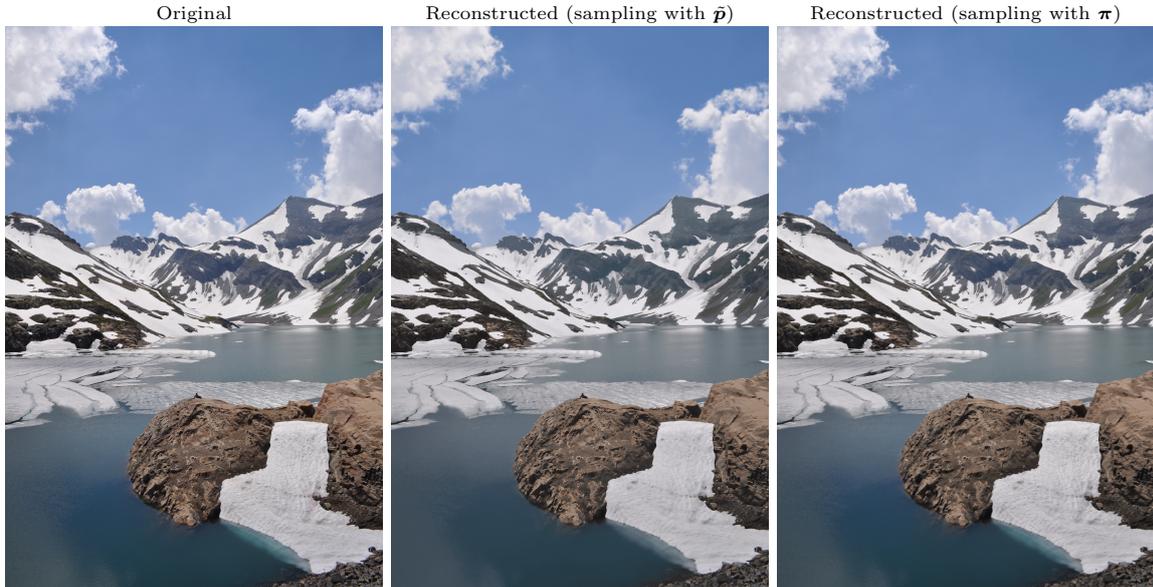


Figure 8. *From left to right: original image; reconstructed image from the measurements obtained with $\tilde{\mathbf{p}}$ (the reconstruction SNR is 27.76 dB); reconstructed image from the measurements obtained with $\boldsymbol{\pi}$ (the reconstruction SNR is 27.10 dB).*

region by copying a single representative patch. In practice this is done via the Laplacian matrix, which encodes the similarities between the patches, by solving (9).

We reconstruct the image by solving (9) for each column of \mathbf{Y} with $\gamma = 1$ and $g(\mathbf{L}) = \mathbf{L}$. Reconstructing the image takes about 3 minutes by solving (10) using the `mldivide` function of Matlab. We show the reconstructed image in Fig. 8. One can notice that we obtain a very accurate reconstruction of the original image. The SNR between the original and the reconstructed images is 27.76 dB. As a comparison, we present in Fig. 8 the reconstructed image from, again, $m = 28622 \approx 0.15n$ measurements but obtained using the uniform sampling distribution. The effective sampling ratio in this case is about 14%. The associated sampled image is presented in Fig. 7(d). The SNR between the original and the reconstructed images is 27.10 dB. The estimated optimal sampling distribution $\tilde{\mathbf{p}}$ allows us to attain a better image quality with an effective sampling ratio almost twice smaller.

6. Conclusion and perspectives

We proposed two efficient sampling procedures for k -bandlimited signals defined on the nodes of a graph \mathcal{G} . The performance of these sampling techniques is governed by the graph weighted coherence, which characterises the interaction between the sampling distribution and the localisation of the first k Fourier modes over the nodes of \mathcal{G} . For regular graph with non-localised Fourier modes and a uniform sampling distribution, we proved that $O(k \log k)$ samples are sufficient to embed the set of k -bandlimited signals. For arbitrary graphs, uniform sampling might perform very poorly. In such cases, we proved that it is always possible to adapt the sampling distribution to the structure of the graph and reach optimal sampling conditions. We designed an algorithm to estimate the optimal sampling distribution rapidly. Finally, we proposed an efficient decoder that provides accurate and stable reconstruction of k -bandlimited signals from their samples.

We believe that the sampling method developed in this work can be used to speed up computations in multiple applications using graph models. Let us take the example of the fast robust PCA method proposed in [62]. In this work, the authors consider the case where one has access to two graphs \mathcal{G}_1 and \mathcal{G}_2 that respectively model the similarities between the rows and the columns of a matrix \mathbf{X} . In this context, they propose an optimisation technique that provides a low-rank approximation of \mathbf{X} . We

denote this low-rank approximation by \mathbf{X}^* . The intuition is that the left singular vectors and the right singular vectors of \mathbf{X}^* live respectively in the span of the first eigenvectors of \mathbf{L}_1 and \mathbf{L}_2 , the Laplacians associated to \mathcal{G}_1 and \mathcal{G}_2 . Therefore, the singular vectors of \mathbf{X}^* can be drastically subsampled using our sampling method. The low-rank matrix \mathbf{X}^* can be reconstructed from a subset of its rows and columns. Instead of estimating \mathbf{X}^* from the entire matrix \mathbf{X} , one could thus first reduce the dimension of the problem by selecting a small subset of the rows and columns of \mathbf{X} .

In semi-supervised learning, a small subset of nodes are labeled and the goal is to infer the label of all nodes. Advances in sampling of graph signals give insight on which nodes should be preferentially observed to infer the labels on the complete graphs. Similarly, in spectral graph clustering, cluster assignments are well approximated by k -bandlimited signals and can therefore be heavily subsampled. This leaves the possibility to initially cluster a small subset of the nodes and infer the clustering solution on the complete graphs afterwards, as we have proposed in [63].

Sensor networks provide other applications of our sampling methods. Indeed, if signals measured by a network of sensors are smooth, one can deduce *beforehand* from the structure of the network which sensors to sample in priority in an active sampling strategy, using the optimal or estimated sampling distribution.

Appendix A - Proof of the theorems in Section 2

We start with the proof of Theorem 2.2. For this proof, we need the following result obtained by Tropp in [64].

Lemma A.1 (Theorem 1.1, [64]). *Consider a finite sequence $\{\mathbf{X}_i\}$ of independent, random, self-adjoint, positive semi-definite matrices of dimension $d \times d$. Assume that each random matrix satisfies*

$$\lambda_{\max}(\mathbf{X}_i) \leq R \quad \text{almost surely.}$$

Define

$$\mu_{\min} := \lambda_{\min} \left(\sum_i \mathbb{E} \mathbf{X}_i \right) \quad \text{and} \quad \mu_{\max} := \lambda_{\max} \left(\sum_i \mathbb{E} \mathbf{X}_i \right).$$

Then

$$\begin{aligned} \mathbb{P} \left\{ \lambda_{\min} \left(\sum_i \mathbf{X}_i \right) \leq (1 - \delta) \mu_{\min} \right\} &\leq d \cdot \left[\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right]^{\mu_{\min}/R} \quad \text{for } \delta \in [0, 1], \quad \text{and} \\ \mathbb{P} \left\{ \lambda_{\max} \left(\sum_i \mathbf{X}_i \right) \geq (1 + \delta) \mu_{\max} \right\} &\leq d \cdot \left[\frac{e^{\delta}}{(1 + \delta)^{1+\delta}} \right]^{\mu_{\max}/R} \quad \text{for } \delta \geq 0. \end{aligned}$$

We also need the following facts. For all $\delta \in [0, 1]$, we have

$$\left[\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right]^{\mu_{\min}/R} \leq \exp \left(-\frac{\delta^2 \mu_{\min}}{3R} \right) \quad \text{and} \quad \left[\frac{e^{\delta}}{(1 + \delta)^{1+\delta}} \right]^{\mu_{\max}/R} \leq \exp \left(-\frac{\delta^2 \mu_{\max}}{3R} \right).$$

Proof of Theorem 2.2. As the i^{th} row-vector of $\mathbf{M}\mathbf{P}^{-1/2}\mathbf{U}_k$ is $\boldsymbol{\delta}_{\omega_i}^{\top} \mathbf{U}_k / \sqrt{p_{\omega_i}}$, we have

$$\frac{1}{m} \mathbf{U}_k^{\top} \mathbf{P}^{-1/2} \mathbf{M}^{\top} \mathbf{M} \mathbf{P}^{-1/2} \mathbf{U}_k = \sum_{i=1}^m \frac{(\mathbf{U}_k^{\top} \boldsymbol{\delta}_{\omega_i}) (\boldsymbol{\delta}_{\omega_i}^{\top} \mathbf{U}_k)}{m p_{\omega_i}}.$$

Let us define

$$\mathbf{X}_i := \frac{1}{m p_{\omega_i}} \mathbf{U}_k^{\top} \boldsymbol{\delta}_{\omega_i} \boldsymbol{\delta}_{\omega_i}^{\top} \mathbf{U}_k,$$

and

$$\mathbf{X} := \sum_{i=1}^m \mathbf{X}_i = m^{-1} \mathbf{U}_k^{\top} \mathbf{P}^{-1/2} \mathbf{M}^{\top} \mathbf{M} \mathbf{P}^{-1/2} \mathbf{U}_k.$$

The matrix \mathbf{X} is thus a sum of m of independent, random, self-adjoint, positive semi-definite matrices. We are in the setting of Lemma A.1. We continue by computing $\mathbb{E} \mathbf{X}_i$ and $\lambda_{\max}(\mathbf{X}_i)$.

The expected value of each \mathbf{X}_i is

$$\mathbb{E} \mathbf{X}_i = \mathbb{E} \left[\frac{(\mathbf{U}_k^\top \boldsymbol{\delta}_{\omega_i}) (\boldsymbol{\delta}_{\omega_i}^\top \mathbf{U}_k)}{m \mathbf{p}_{\omega_i}} \right] = \frac{1}{m} \mathbf{U}_k^\top \left(\sum_{i=1}^n \mathbf{p}_i \frac{\boldsymbol{\delta}_i \boldsymbol{\delta}_i^\top}{\mathbf{p}_i} \right) \mathbf{U}_k = \frac{1}{m} \mathbf{U}_k^\top \mathbf{U}_k = \frac{1}{m} \mathbf{I}$$

where $\mathbf{I} \in \mathbb{R}^{k \times k}$ is the identity matrix. Therefore,

$$\mu_{\min} := \lambda_{\min} \left(\sum_i \mathbb{E} \mathbf{X}_i \right) = 1 \quad \text{and} \quad \mu_{\max} := \lambda_{\max} \left(\sum_i \mathbb{E} \mathbf{X}_i \right) = 1.$$

Furthermore, for all $i = 1, \dots, n$, we have

$$\lambda_{\max}(\mathbf{X}_i) = \|\mathbf{X}_i\|_2 \leq \max_{1 \leq j \leq n} \left\| \frac{(\mathbf{U}_k^\top \boldsymbol{\delta}_j) (\boldsymbol{\delta}_j^\top \mathbf{U}_k)}{m \mathbf{p}_j} \right\|_2 = \frac{1}{m} \max_{1 \leq j \leq n} \left\{ \frac{\|\mathbf{U}_k^\top \boldsymbol{\delta}_j\|_2^2}{\mathbf{p}_j} \right\} = \frac{(\nu_{\mathbf{p}}^k)^2}{m}.$$

Lemma A.1 yields, for any $\delta \in (0, 1)$,

$$\begin{aligned} \mathbb{P} \{ \lambda_{\min}(\mathbf{X}) \leq (1 - \delta) \} &\leq k \cdot \left[\frac{e^{-\delta}}{(1 - \delta)^{1 - \delta}} \right]^{m / (\nu_{\mathbf{p}}^k)^2} \leq k \exp \left(-\frac{\delta^2 m}{3 (\nu_{\mathbf{p}}^k)^2} \right) \quad \text{and} \\ \mathbb{P} \{ \lambda_{\max}(\mathbf{X}) \geq (1 + \delta) \} &\leq k \cdot \left[\frac{e^{\delta}}{(1 + \delta)^{1 + \delta}} \right]^{m / (\nu_{\mathbf{p}}^k)^2} \leq k \exp \left(-\frac{\delta^2 m}{3 (\nu_{\mathbf{p}}^k)^2} \right). \end{aligned}$$

Therefore, for any $\delta \in (0, 1)$, we have, with probability at least $1 - \epsilon$,

$$(15) \quad 1 - \delta \leq \lambda_{\min}(\mathbf{X}) \quad \text{and} \quad \lambda_{\max}(\mathbf{X}) \leq 1 + \delta$$

provided that

$$m \geq \frac{3}{\delta^2} (\nu_{\mathbf{p}}^k)^2 \log \left(\frac{2k}{\epsilon} \right).$$

Noticing that (15) implies that

$$(1 + \delta) \|\boldsymbol{\alpha}\|_2^2 \leq \left\| \mathbf{M} \mathbf{P}^{-1/2} \mathbf{U}_k \boldsymbol{\alpha} \right\|_2^2 \leq (1 + \delta) \|\boldsymbol{\alpha}\|_2^2,$$

for all $\boldsymbol{\alpha} \in \mathbb{R}^k$, which is equivalent to

$$(1 + \delta) \|\mathbf{x}\|_2^2 \leq \left\| \mathbf{M} \mathbf{P}^{-1/2} \mathbf{x} \right\|_2^2 \leq (1 + \delta) \|\mathbf{x}\|_2^2,$$

for all $\mathbf{x} \in \text{span}(\mathbf{U}_k)$, terminates the proof. \square

The proof of Theorem 2.4 is based on the following results, also obtained by Tropp.

Lemma A.2 (Theorem 2.2, [65]). *Let \mathcal{X} be a finite set of positive-semidefinite matrices of dimension $d \times d$, and suppose that*

$$\max_{\mathbf{X} \in \mathcal{X}} \lambda_{\max}(\mathbf{X}) \leq R.$$

Sample $\{\mathbf{X}_1, \dots, \mathbf{X}_l\}$ uniformly at random from \mathcal{X} without replacement. Compute

$$\mu_{\min} := l \cdot \lambda_{\min}(\mathbb{E} \mathbf{X}_1) \quad \text{and} \quad \mu_{\max} := l \cdot \lambda_{\max}(\mathbb{E} \mathbf{X}_1).$$

Then

$$\begin{aligned} \mathbb{P} \left\{ \lambda_{\min} \left(\sum_i \mathbf{X}_i \right) \leq (1 - \delta) \mu_{\min} \right\} &\leq d \cdot \left[\frac{e^{-\delta}}{(1 - \delta)^{1 - \delta}} \right]^{\mu_{\min}/R} \quad \text{for } \delta \in [0, 1], \quad \text{and} \\ \mathbb{P} \left\{ \lambda_{\max} \left(\sum_i \mathbf{X}_i \right) \geq (1 + \delta) \mu_{\max} \right\} &\leq d \cdot \left[\frac{e^{\delta}}{(1 + \delta)^{1 + \delta}} \right]^{\mu_{\max}/R} \quad \text{for } \delta \geq 0. \end{aligned}$$

Using Lemma A.1, one can notice that the above probability bounds would be identical if the matrices $\{X_1, \dots, X_l\}$ were sampled uniformly at random from \mathcal{X} *with* replacement. It is thus not necessary to detail the complete proof which is entirely similar to the one of Theorem 2.2, at the exception of the sampling procedure.

Appendix B - Proof of the theorems in Section 3

Proof of Theorem 3.1. We recall that \mathbf{x}^* is a solution to (5). By optimality of \mathbf{x}^* , we have

$$\left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} \mathbf{x}^* - \mathbf{P}_\Omega^{-1/2} \mathbf{y} \right\|_2 \leq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} \mathbf{z} - \mathbf{P}_\Omega^{-1/2} \mathbf{y} \right\|_2$$

for any $\mathbf{z} \in \text{span}(\mathbf{U}_k)$. In particular for $\mathbf{z} = \mathbf{x}$, we obtain

$$\left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} \mathbf{x}^* - \mathbf{P}_\Omega^{-1/2} \mathbf{y} \right\|_2 \leq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} \mathbf{x} - \mathbf{P}_\Omega^{-1/2} \mathbf{y} \right\|_2,$$

which yields

$$(16) \quad \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} \mathbf{x}^* - \mathbf{P}_\Omega^{-1/2} \mathbf{M} \mathbf{x} - \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 \leq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2.$$

Then, the triangle inequality and (2) yields

$$(17) \quad \begin{aligned} \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} (\mathbf{x}^* - \mathbf{x}) - \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 &\geq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} (\mathbf{x}^* - \mathbf{x}) \right\|_2 - \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 \\ &= \left\| \mathbf{M} \mathbf{P}^{-1/2} (\mathbf{x}^* - \mathbf{x}) \right\|_2 - \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 \\ &\geq \sqrt{m(1-\delta)} \|\mathbf{x}^* - \mathbf{x}\|_2 - \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2. \end{aligned}$$

In the second step, we used the fact that $\mathbf{M} \mathbf{P}^{-1/2} = \mathbf{P}_\Omega^{-1/2} \mathbf{M}$. Combining (16) and (17) directly yields (6), the first bound in Theorem 3.1.

To prove the second bound, let us choose $\mathbf{n}_0 = \mathbf{M} \mathbf{z}_0$ with $\mathbf{z}_0 \in \text{span}(\mathbf{U}_k)$. Therefore, $\mathbf{y} = \mathbf{M}(\mathbf{x} + \mathbf{z}_0)$ and $\mathbf{x}^* = \mathbf{x} + \mathbf{z}_0$ is an obvious solution to (5) in this case. To finish the proof, we use (2) which yields

$$\begin{aligned} \|\mathbf{x}^* - \mathbf{x}\|_2 = \|\mathbf{z}_0\|_2 &\geq \frac{1}{\sqrt{m(1+\delta)}} \left\| \mathbf{M} \mathbf{P}^{-1/2} \mathbf{z}_0 \right\|_2 = \frac{1}{\sqrt{m(1+\delta)}} \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M} \mathbf{z}_0 \right\|_2 \\ &= \frac{1}{\sqrt{m(1+\delta)}} \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n}_0 \right\|_2. \end{aligned}$$

□

Proof of Theorem 3.2. As \mathbf{x}^* is a solution to (9), we have

$$(18) \quad \left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M} \mathbf{x}^* - \mathbf{y}) \right\|_2^2 + \gamma (\mathbf{x}^*)^\top g(\mathbf{L}) \mathbf{x}^* \leq \left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M} \mathbf{z} - \mathbf{y}) \right\|_2^2 + \gamma \mathbf{z}^\top g(\mathbf{L}) \mathbf{z},$$

for all $\mathbf{z} \in \mathbb{R}^n$. We also have $\mathbf{x}^* = \boldsymbol{\alpha}^* + \boldsymbol{\beta}^*$ with $\boldsymbol{\alpha}^* \in \text{span}(\mathbf{U}_k)$ and $\boldsymbol{\beta}^* \in \text{span}(\bar{\mathbf{U}}_k)$. Let us define the matrix

$$\bar{\mathbf{U}}_k := (\mathbf{u}_{k+1}, \dots, \mathbf{u}_n) \in \mathbb{R}^{n \times (n-k)}.$$

Choosing $\mathbf{z} = \mathbf{x}$ in (18) and using the facts that $\bar{\mathbf{U}}_k^\top \boldsymbol{\alpha}^* = \mathbf{0}$, $\mathbf{U}_k^\top \boldsymbol{\beta}^* = \mathbf{0}$, $\bar{\mathbf{U}}_k^\top \mathbf{x} = \mathbf{0}$, and that $g(\mathbf{L}) = \mathbf{U} g(\mathbf{L}) \mathbf{U}^\top$, we obtain

$$\begin{aligned} \left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M} \mathbf{x}^* - \mathbf{y}) \right\|_2^2 + \gamma (\mathbf{U}_k^\top \boldsymbol{\alpha}^*)^\top \mathbf{G}_k (\mathbf{U}_k^\top \boldsymbol{\alpha}^*) + \gamma (\bar{\mathbf{U}}_k^\top \boldsymbol{\beta}^*)^\top \bar{\mathbf{G}}_k (\bar{\mathbf{U}}_k^\top \boldsymbol{\beta}^*) \\ \leq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2^2 + \gamma (\mathbf{U}_k^\top \mathbf{x})^\top \mathbf{G}_k (\mathbf{U}_k^\top \mathbf{x}), \end{aligned}$$

where

$$\mathbf{G}_k := \text{diag}(g(\boldsymbol{\lambda}_1), \dots, g(\boldsymbol{\lambda}_k)) \in \mathbb{R}^{k \times k} \text{ and } \bar{\mathbf{G}}_k := \text{diag}(g(\boldsymbol{\lambda}_{k+1}), \dots, g(\boldsymbol{\lambda}_n)) \in \mathbb{R}^{(n-k) \times (n-k)}.$$

We deduce that

$$\left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M}\mathbf{x}^* - \mathbf{y}) \right\|_2^2 + \gamma g(\boldsymbol{\lambda}_{k+1}) \|\boldsymbol{\beta}^*\|_2^2 \leq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2^2 + \gamma g(\boldsymbol{\lambda}_k) \|\mathbf{x}\|_2^2,$$

where we used the fact that $\|\bar{\mathbf{U}}_k^\top \boldsymbol{\beta}^*\|_2 = \|\boldsymbol{\beta}^*\|_2$ and $\|\mathbf{U}_k^\top \mathbf{x}\|_2 = \|\mathbf{x}\|_2$. As the left hand side of the last inequality is a sum of two positive quantities, we also have

$$(19) \quad \left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M}\mathbf{x}^* - \mathbf{y}) \right\|_2 \leq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 + \sqrt{\gamma g(\boldsymbol{\lambda}_k)} \|\mathbf{x}\|_2 \text{ and}$$

$$(20) \quad \sqrt{\gamma g(\boldsymbol{\lambda}_{k+1})} \|\boldsymbol{\beta}^*\|_2 \leq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 + \sqrt{\gamma g(\boldsymbol{\lambda}_k)} \|\mathbf{x}\|_2.$$

Inequality (20) proves (12), the second inequality in Theorem 3.2. It remains to prove (11). To prove this inequality, we continue by using (2), which yields

$$(21) \quad \begin{aligned} \left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M}\mathbf{x}^* - \mathbf{y}) \right\|_2 &= \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M}(\boldsymbol{\alpha}^* - \mathbf{x}) + \mathbf{P}_\Omega^{-1/2} \mathbf{n} + \mathbf{P}_\Omega^{-1/2} \mathbf{M}\boldsymbol{\beta}^* \right\|_2 \\ &\geq \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M}(\boldsymbol{\alpha}^* - \mathbf{x}) \right\|_2 - \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 - \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{M}\boldsymbol{\beta}^* \right\|_2 \\ &= \left\| \mathbf{M}\mathbf{P}_\Omega^{-1/2} (\boldsymbol{\alpha}^* - \mathbf{x}) \right\|_2 - \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 - \left\| \mathbf{M}\mathbf{P}_\Omega^{-1/2} \boldsymbol{\beta}^* \right\|_2 \\ &\geq \sqrt{m(1-\delta)} \|\boldsymbol{\alpha}^* - \mathbf{x}\|_2 - \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 - M_{\max} \|\boldsymbol{\beta}^*\|_2. \end{aligned}$$

Finally, combining (19), (20) and (21) gives

$$\begin{aligned} \|\boldsymbol{\alpha}^* - \mathbf{x}\|_2 &\leq \frac{1}{\sqrt{m(1-\delta)}} \left(\left\| \mathbf{P}_\Omega^{-1/2} (\mathbf{M}\mathbf{x}^* - \mathbf{y}) \right\|_2 + \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 + M_{\max} \|\boldsymbol{\beta}^*\|_2 \right) \\ &\leq \frac{1}{\sqrt{m(1-\delta)}} \left(\left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 + \sqrt{\gamma g(\boldsymbol{\lambda}_k)} \|\mathbf{x}\|_2 \right. \\ &\quad \left. + \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 \right. \\ &\quad \left. + \frac{M_{\max} \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2}{\sqrt{\gamma g(\boldsymbol{\lambda}_{k+1})}} + M_{\max} \sqrt{\frac{g(\boldsymbol{\lambda}_k)}{g(\boldsymbol{\lambda}_{k+1})}} \|\mathbf{x}\|_2 \right) \\ &\leq \frac{1}{\sqrt{m(1-\delta)}} \left(2 + \frac{M_{\max}}{\sqrt{\gamma g(\boldsymbol{\lambda}_{k+1})}} \right) \left\| \mathbf{P}_\Omega^{-1/2} \mathbf{n} \right\|_2 \\ &\quad + \frac{1}{\sqrt{m(1-\delta)}} \left(M_{\max} \sqrt{\frac{g(\boldsymbol{\lambda}_k)}{g(\boldsymbol{\lambda}_{k+1})}} + \sqrt{\gamma g(\boldsymbol{\lambda}_k)} \right) \|\mathbf{x}\|_2. \end{aligned}$$

This terminates the proof. \square

Appendix C - Proof of the theorem in Section 4

We use the classical technique to prove the Johnson-Lindenstrauss lemma (see, *e.g.*, [66]).

Proof. Each filtered signal $\mathbf{r}_{\hat{c}_\lambda}^l$, $l \in \{1, \dots, L\}$, satisfies

$$\mathbf{r}_{\hat{c}_\lambda}^l = \mathbf{U} \mathbf{C}_\lambda \mathbf{U}^\top \mathbf{r}^l,$$

where $\mathbf{C}_\lambda := \text{diag}(\hat{c}_\lambda(\boldsymbol{\lambda}_1), \dots, \hat{c}_\lambda(\boldsymbol{\lambda}_n))$. Let i be fixed for the moment. We have

$$\sum_{l=1}^L (\mathbf{r}_{\hat{c}_\lambda}^l)_i^2 = \sum_{l=1}^L (\delta_i^\top \mathbf{U} \mathbf{C}_\lambda \mathbf{U}^\top \mathbf{r}^l)^2,$$

The expected value of this sum is $\|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2$. Indeed,

$$\begin{aligned} \mathbb{E} \left[\sum_{l=1}^L (\mathbf{r}_{\hat{c}_\lambda}^l)_i^2 \right] &= \sum_{l=1}^L \boldsymbol{\delta}_i^\top \mathbf{U} \mathbf{C}_\lambda \mathbf{U}^\top \mathbb{E} [\mathbf{r}^l (\mathbf{r}^l)^\top] \mathbf{U} \mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i = L^{-1} \sum_{l=1}^L \boldsymbol{\delta}_i^\top \mathbf{U} \mathbf{C}_\lambda \mathbf{U}^\top \mathbf{U} \mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i \\ &= \boldsymbol{\delta}_i^\top \mathbf{U} \mathbf{C}_\lambda^2 \mathbf{U}^\top \boldsymbol{\delta}_i = \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2. \end{aligned}$$

Let us define

$$X_i := \sum_{l=1}^L \left[(\mathbf{r}_{\hat{c}_\lambda}^l)_i^2 - L^{-1} \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2 \right].$$

This is a sum of L independent centered random variables. Furthermore, as each \mathbf{r}^l is a zero-mean Gaussian random vector with covariance matrix $L^{-1} \mathbf{I}$, the variables $(\mathbf{r}_{\hat{c}_\lambda}^l)_i$ are subgaussian with subgaussian bounded by $C L^{-1/2} \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2$, where $C \geq 1$ is an absolute constant. We let the reader refer to, *e.g.*, [67] for more information on the definition and properties of subgaussian random variables. Using Lemma 5.14 and Remark 5.18 in [67], one can prove that each summand of X is a centered subexponential random variable with subexponential norm bounded by $4C^2 L^{-1} \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2$. Corollary 5.17 in [67] shows that there exists an absolute constant $c > 0$ such that

$$\mathbb{P}(|X_i| \geq tL) \leq 2 \exp \left(-\frac{c L t^2}{16C^4 L^{-2} \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^4} \right)$$

for all $t \in (0, 4C^2 L^{-1} \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2)$, or, equivalently, that

$$\mathbb{P}(|X_i| \geq \delta \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2) \leq 2 \exp \left(-\frac{c L \delta^2}{16C^4} \right),$$

for all $\delta \in (0, 4C^2)$.

Then, using the union bound, we obtain

$$\mathbb{P} \left(\max_{i \in \{1, \dots, n\}} |X_i| \geq \delta \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2 \right) \leq 2n \exp \left(-\frac{c L \delta^2}{16C^4} \right).$$

This proves that, with probability at least $1 - \epsilon$,

$$(1 - \delta) \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2 \leq \sum_{l=1}^L (\mathbf{r}_{\hat{c}_\lambda}^l)_i^2 \leq (1 + \delta) \|\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i\|_2^2,$$

for all $i \in \{1, \dots, n\}$, provided that

$$L \geq \frac{16C^4}{c \delta^2} \log \left(\frac{2n}{\epsilon} \right).$$

To finish the the proof, one just needs to remark that

$$\mathbf{C}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i = \mathbf{U}_{j^*}^\top \boldsymbol{\delta}_i + \mathbf{E}_\lambda \mathbf{U}^\top \boldsymbol{\delta}_i,$$

by definition of \hat{c}_λ (see (13)) and use the triangle inequality. \square

References

- [1] M. Newman, *Networks: an introduction*. Oxford University Press, 2010.
- [2] D. Shuman, S. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *Signal Processing Magazine, IEEE*, vol. 30, no. 3, pp. 83–98, May 2013.
- [3] A. Sandryhaila and J. Moura, "Big data analysis with signal processing on graphs: Representation and processing of massive data sets with irregular structure," *Signal Processing Magazine, IEEE*, vol. 31, no. 5, pp. 80–90, Sept 2014.
- [4] M. Unser, "Sampling-50 years after shannon," *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569–587, April 2000.
- [5] C. Shannon, "Communication in the presence of noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, Jan 1949.
- [6] J. D. McEwen and Y. Wiaux, "A novel sampling theorem on the sphere," *IEEE Trans. Signal Process.*, vol. 59, no. 12, 2011.

- [7] K. Gröchenig, “Reconstruction algorithms in irregular sampling,” *Mathematics of Computation*, vol. 59, no. 199, pp. 181–194, 1992.
- [8] J. J. Benedetto, “Irregular sampling and frames,” *wavelets: A Tutorial in Theory and Applications*, vol. 2, pp. 445–507, 1992.
- [9] E. J. Candès *et al.*, “Compressive sampling,” in *Proceedings of the international congress of mathematicians*, vol. 3. Madrid, Spain, 2006, pp. 1433–1452.
- [10] S. Chen, R. Varma, A. Sandryhaila, and J. Kovacevic, “Discrete signal processing on graphs: Sampling theory,” *Signal Processing, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.
- [11] A. Anis, A. Gadde, and A. Ortega, “Towards a sampling theorem for signals on arbitrary graphs,” in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 3864–3868.
- [12] S. Narang and A. Ortega, “Compact support biorthogonal wavelet filterbanks for arbitrary undirected graphs,” *Signal Processing, IEEE Transactions on*, vol. 61, no. 19, pp. 4673–4685, Oct 2013.
- [13] I. Pesenson, “Sampling in Paley-Wiener spaces on combinatorial graphs,” *Transactions of the American Mathematical Society*, vol. 360, no. 10, pp. 5603–5627, 2008.
- [14] I. Z. Pesenson and M. Z. Pesenson, “Sampling, Filtering and Sparse Approximations on Combinatorial Graphs,” *J. Fourier Anal. Appl.*, vol. 16, no. 6, pp. 921–942, 2010.
- [15] A. Anis, A. E. Gamal, S. Avestimehr, and A. Ortega, “Asymptotic justification of bandlimited interpolation of graph signals for semi-supervised learning,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2015, pp. 5461–5465.
- [16] S. Chen, A. Sandryhaila, and J. Kovacevic, “Sampling theory for graph signals,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, April 2015, pp. 3392–3396.
- [17] A. Anis, A. Gadde, and A. Ortega, “Efficient sampling set selection for bandlimited graph signals using graph spectral proxies,” *IEEE Transactions on Signal Processing*, vol. PP, no. 99, pp. 1–1, 2016.
- [18] H. Nguyen and M. Do, “Downsampling of signals on graphs via maximum spanning trees,” *Signal Processing, IEEE Transactions on*, vol. 63, no. 1, pp. 182–191, Jan 2015.
- [19] D. I. Shuman, M. J. Faraji, and P. Vandergheynst, “A multiscale pyramid transform for graph signals,” *IEEE Transactions on Signal Processing*, vol. 64, no. 8, pp. 2119–2134, April 2016.
- [20] J. Irion and S. Naoki, “Applied and computational harmonic analysis on graphs and networks,” in *Proc. SPIE Conf. WAVELET XVI*, vol. 9597, 2015.
- [21] N. Tremblay and P. Borgnat, “Subgraph-based filterbanks for graph signals,” *IEEE Transactions on Signal Processing*, vol. PP, no. 99, pp. 1–1, 2016.
- [22] A. Agaskar and Y. Lu, “A spectral graph uncertainty principle,” *Information Theory, IEEE Transactions on*, vol. 59, no. 7, pp. 4338–4356, July 2013.
- [23] M. Rabbat and V. Gripon, “Towards a spectral characterization of signals supported on small-world networks,” in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 4793–4797.
- [24] Y. Nakatsukasa, N. Saito, and E. Woei, “Mysteries around the graph laplacian eigenvalue 4,” *Linear Algebra and its Applications*, vol. 438, no. 8, pp. 3231 – 3246, 2013.
- [25] G. Puy, P. Vandergheynst, and Y. Wiaux, “On variable density compressive sampling,” *IEEE Signal Process. Lett.*, vol. 18, no. 10, pp. 595–598, 2011.
- [26] F. Krahermer and R. Ward, “Stable and robust sampling strategies for compressive imaging,” *IEEE Trans. Image Process.*, vol. 23, no. 2, pp. 612–622, 2013.
- [27] B. Adcock, A. C. Hansen, C. Poon, and B. Roman, “Breaking the coherence barrier: a new theory for compressed sensing,” *arXiv:1302.0561*, 2013.
- [28] S. Chen, R. Varma, A. Singh, and J. Kovacevi, “Signal recovery on graphs: Random versus experimentally designed sampling,” in *Sampling Theory and Applications (SampTA), 2015 International Conference on*, May 2015, pp. 337–341.
- [29] F. Chung, *Spectral graph theory*. Amer Mathematical Society, 1997, no. 92.
- [30] E. Candes and J. Romberg, “Sparsity and incoherence in compressive sampling,” *Inverse Problems*, vol. 23, no. 3, pp. 969–985, 2007.
- [31] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, ser. Applied and Numerical Harmonic Analysis. Springer New York, 2013.
- [32] M. Lustig, D. Donoho, and J. M. Pauly, “Sparse MRI: The application of compressed sensing for rapid MR imaging,” *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [33] N. Chauffert, P. Ciuciu, and P. Weiss, “Variable density compressed sensing in mri. theoretical vs heuristic sampling strategies,” in *Proc. IEEE Int. Symp. on Biomedical Imaging*, 2013, pp. 298 – 301.
- [34] C. Boyer, J. Bigot, and P. Weiss, “Compressed sensing with structured sparsity and structured acquisition,” *arXiv:1505.01619*, 2015.
- [35] J. Bigot, C. Boyer, and P. Weiss, “An analysis of block sampling strategies in compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 62, no. 4, pp. 2125 – 2139, 2016.
- [36] P. Drineas, M. M.-Ismail, M. W. Mahoney, and D. P. Woodruff, “Fast approximation of matrix coherence and statistical leverage,” *J. Mach. Learn. Res.*, vol. 13, p. 34753506, 2012.

- [37] P. Drineas, M. W. Mahoney, and S. Muthukrishnan, “Sampling algorithms for l2 regression and applications,” in *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm (SODA’06)*, 2006, pp. 1127–1136.
- [38] M. W. Mahoney, “Randomized algorithms for matrices and data,” *Foundation and Trends in Machine Learning*, vol. 3, no. 2, pp. 123–224, 2011.
- [39] S. Chen, R. Varma, A. Singh, and J. Kovacevic, “A statistical perspective of sampling scores for linear regression,” 2016, submitted.
- [40] P. Drineas, M. W. Mahoney, and S. Muthukrishnan, “Relative-error *cur* matrix decompositions,” *SIAM J. Matrix Anal. & Appl.*, vol. 30, no. 2, pp. 844–881, 2008.
- [41] A. Gittens and M. Mahoney, “Revisiting the nystrom method for improved large-scale machine learning,” *J. Mach. Learn. Res.*, vol. 28, no. 3, pp. 567–575, 2013.
- [42] S. Sun, J. Zhao, and J. Zhu, “A review of nystrom methods for large-scale machine learning,” *Information Fusion*, vol. 26, pp. 35–48, 2015.
- [43] D. K. Hammond, P. Vandergheynst, and R. Gribonval, “Wavelets on graphs via spectral graph theory,” *Appl. Comput. Harmon. Anal.*, vol. 30, no. 2, pp. 129–150, 2011.
- [44] O. Chapelle, B. Schölkopf, and A. Zien, *Semi-Supervised Learning*, 1st ed. The MIT Press, 2010.
- [45] X. Zhu, Z. Ghahramani, and J. D. Lafferty, “Semi-supervised learning using gaussian fields and harmonic functions,” in *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21–24, Washington, DC, USA, 2003*, pp. 912–919.
- [46] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf, “Learning with local and global consistency,” in *Advances in Neural Information Processing Systems 16*, S. Thrun, L. K. Saul, and B. Schölkopf, Eds. MIT Press, 2004, pp. 321–328.
- [47] M. Belkin, P. Niyogi, and V. Sindhwani, “Manifold regularization: A geometric framework for learning from labeled and unlabeled examples,” *J. Mach. Learn. Res.*, vol. 7, pp. 2399–2434, Dec. 2006.
- [48] Y. Bengio, O. Delalleau, and N. L. Roux, *Label Propagation and Quadratic Criterion*. MIT Press, 2006, p. 193216.
- [49] Y.-X. Wang, J. Sharpnack, A. Smola, and R. J. Tibshirani, “Trend filtering on graphs,” in *International Conference on Artificial Intelligence and Statistics*, 2015, pp. 1042–1050.
- [50] M. Belkin and P. Niyogi, “Using manifold structure for partially labeled classification,” in *Advances in Neural Information Processing Systems 15*, S. Becker, S. Thrun, and K. Obermayer, Eds. MIT Press, 2003, pp. 953–960.
- [51] R. Fergus, Y. Weiss, and A. Torralba, “Semi-supervised learning in gigantic image collections,” in *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, Eds. Curran Associates, Inc., 2009, pp. 522–530.
- [52] B. Settles, *Active learning*. Morgan & Claypool Publishers, 2012, vol. 6, no. 1.
- [53] Y. Fu, X. Zhu, and B. Li, “A survey on instance selection for active learning,” *Knowledge and Information Systems*, vol. 35, no. 2, pp. 249–283, 2012.
- [54] Q. Gu and J. Han, “Towards active learning on graphs: An error bound minimization approach,” in *2012 IEEE 12th International Conference on Data Mining*, Dec 2012, pp. 882–887.
- [55] M. Bilgic and L. Getoor, “Active inference for collective classification,” *AAAI Conference on Artificial Intelligence*, 2010.
- [56] M. Ji and J. Han, “A variance minimization criterion to active learning on graphs,” in *International Conference on Artificial Intelligence and Statistics*, 2012, pp. 556–564.
- [57] Y. Ma, R. Garnett, and J. Schneider, “ σ -optimality for active learning on gaussian random fields,” in *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2013, pp. 2751–2759.
- [58] C. Zhang, D. Florêncio, and P. A. Chou, “Graph signal processing—a probabilistic framework,” Microsoft Research, Tech. Rep., April 2015. [Online]. Available: <http://research.microsoft.com/apps/pubs/default.aspx?id=243326>
- [59] A. Gadde and A. Ortega, “A probabilistic interpretation of sampling theory of graph signals,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2015, pp. 3257–3261.
- [60] E. D. Napoli, E. Polizzo, and Y. Saad, “Efficient estimation of eigenvalue counts in an interval,” *arXiv:1308.4275*, 2013.
- [61] N. Perraudin, J. Paratte, D. Shuman, V. Kalofolias, P. Vandergheynst, and D. K. Hammond, “Gspbox: A toolbox for signal processing on graphs,” *arXiv:1408.5781*, 2014.
- [62] N. Shahid, N. Perraudin, V. Kalofolias, and P. Vandergheynst, “Fast robust PCA on graphs,” *arXiv:1507:08173*, 2015.
- [63] N. Tremblay, G. Puy, R. Gribonval, and P. Vandergheynst, “Compressive spectral clustering,” in *Machine Learning, Proceedings of the Thirty-third International Conference (ICML 2016), June 20–22, New York City, USA, 2016*.
- [64] J. A. Tropp, “User-friendly tail bounds for sums of random matrices,” *Found. Comput. Math.*, vol. 12, no. 4, pp. 389–434, 2012.
- [65] —, “Improved analysis of the subsampled randomized hadamard transform,” *Adv. Adapt. Data Anal.*, vol. 3, no. 01n02, pp. 115–126, 2011.
- [66] R. G. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, “A simple proof of the restricted isometry property for random matrices,” *Constr. Approx.*, vol. 28, no. 3, pp. 253–263, 2008.

- [67] R. Vershynin, *Compressed Sensing, Theory and Applications*. Cambridge University Press, 2012, ch. Introduction to the non-asymptotic analysis of random matrices, pp. 210–268.