Graphon Pooling for Reducing Dimensionality of Signals and Convolutional Operators on Graphs

Alejandro Parada-Mayorga, Zhiyang Wang, and Alejandro Ribeiro

arXiv:2212.08171v1 [cs.LG] 15 Dec 2022

Abstract—In this paper we propose a pooling approach for convolutional information processing on graphs relying on the theory of graphons and limits of dense graph sequences. We present three methods that exploit the induced graphon representation of graphs and graph signals on partitions of $[0, 1]^2$ in the graphon space. As a result we derive low dimensional representations of the convolutional operators, while a dimensionality reduction of the signals is achieved by simple local interpolation of functions in $L^2([0, 1])$. We prove that those low dimensional representations constitute a convergent sequence of graphs and graph signals, respectively. The methods proposed and the theoretical guarantees that we provide show that the reduced graphs and signals inherit spectral-structural properties of the original quantities. We evaluate our approach with a set of numerical experiments performed on graph neural networks (GNNs) that rely on graphon pooling. We observe that graphon pooling performs significantly better than other approaches proposed in the literature when dimensionality reduction ratios between layers are large. We also observe that when graphon pooling is used we have, in general, less overfitting and lower computational cost.

Index Terms—Graphon pooling, convolutional operators on graphs, graphons, graphon signal processing, graphon neural networks, dense graph limits.

I. INTRODUCTION

The problem of pooling entails finding lower dimensional representations of *signals* defined on a given domain and *operators* acting on these signals [2]–[6]. A good pooling technique is such that the result of applying pooled operators to pooled signals yields outputs that are close to pooled versions of the result of applying the original operator to the original signal. In the classical case of signals in time or space processed with convolutional filters or convolutional neural networks (CNNs), pooling is rather straightforward. A pooled signal is a local average and a pooled filter is a sampled filter. Provided that the signal or the filter contain most of their energy in sufficiently low frequencies, the sampling theorem guarantees that processing the pooled signal with the pooled (sampled) filter is a good approximation of the processing of the original signal with the original filter [7, Section 1.7].

In this paper we consider signals supported on graphs – large graphs in particular – and graph convolutional operators in the form of graph filters and graph neural networks (GNNs). In this case, pooling is more involved. To pool signals we can still consider some simple local averages. We can also invoke results from the graph signal processing literature to ascertain

that the pooled signal and the original signal are similar representations if the energy of the signal is concentrated in low graph frequencies [8]–[17]. The challenge is that pooled graph filters are not easy to devise. This is because graph convolutions are polynomials on matrix representations of the graph [18]–[22]. Thus, once we pool information on a subset of nodes it is unclear how to construct a graph linking these nodes to yield graph convolutions that are good approximations of graph convolutions in the original graph [23]–[25].

To perform pooling on graphs, two main approaches stand out: Multi-scale or multi-level clustering [24], [25] and zero padding [23]. In multi-level clustering – also known as graph coarsening –, families of graphs are derived by grouping subsets of nodes in the original graph. Each cluster is associated to a node in the pooled graph and inter-cluster connectivity determines edges in the pooled graph [24], [25]. Although satisfactory results are obtained with this approach – tested on GNNs –, the computational cost is high; a fact that hinders applicability to large graphs. With zero padding the dimensionality of signals and operators is reduced by zeroing specific components of the signal while retaining the original graph [23]. This procedure forces a reduction in the dimension of the signal while inducing a reduction of the effective dimension of the filtering operators. The effectiveness of zero padding depends on the effectiveness of the choice of the set of zeroed nodes. Finding good sets of nodes to zero is computationally expensive and, as is the case of graph coarsening, precludes application to large graphs.

In this paper we consider graphs with large numbers of nodes. We leverage the concept of graphons as graph limits and build on the theory of graphon signal processing [26]–[28] to provide the following contribution:

(C1) We propose low computational cost pooling methods for signals and operators on graphs based on graphon representations (Section IV).

In particular, we derive an operation of pooling by building sequences of graphs and graph signals that converge to a graphon and a graphon signal, respectively. To build these sequences we leverage results about partitions in graphon spaces, performing three simple operations: (i) integration on a regular grid, (ii) integration on an irregular grid, and (iii) random sampling. The graphs obtained by these methods define the reduced operators on the graph, while the reduction of the signal is achieved by simple local interpolation of functions on $L^2([0, 1])$.

However simple, graphon pooling methods can yield good pooled representations. We prove that this is true by providing the following two contributions:

(C2) We consider pooled versions of signals and filters and provide an error bound for the difference between the outcome of processing the pooled signal with the pooled filter and the outcome of processing the original signal with the original filter. These bounds require conditions on graphon filters that are akin to the low frequency conditions that appear in the processing of time signals (Theorems 3 and 4).

(C3) We show that the shift operators obtained by means of graphon pooling based on integration on a regular grid are stable with respect to arbitrary approximations of the graphon (Theorem 5).

Contribution (C3) is important because contribution (C2) requires access to the graphon from which graphs are sampled. Indeed, our results show that graphon pooling methods induce filter perturbations that are bounded by the cut-norm distance between the original graph and its pooled version provided that both are sampled from the same graphon. In practice, graphons must be estimated from graphs. Contribution (C3) implies that graphon estimation errors stay contained when they are mapped to filter perturbation bounds of pooled operators.

The performance of graphon pooling is tested on GNNs, allowing a direct comparison with the graph pooling methods proposed in [23]–[25]. A set of numerical experiments corroborate our findings and show that when the dimensionality reduction of the signals and operators is large, graphon pooling leads to better performance than other approaches (Section V). In particular, the best performance is obtained by the graphon pooling method based on integration on a regular grid. This is partly explained by the stability of that method to approximations of the graphon.

This paper is organized as follows. In Section II we discuss the basics of graph signal processing and GNNs, including the concepts of signals, convolutional operators (filters), and GNN mapping operators. This will provide the scenario where graphon pooling is naturally applied and where it is also going to be tested numerically. Section III contains the basics about graphons, graph limits, graphon neural networks (Gphon-NN), and their connection to graphs and GNNs. In Section IV we introduce the graphon pooling methods proposed in this paper, while in Section V we provide numerical simulations to evaluate the performance of graphon pooling against other pooling approaches. In Section VI we present discussions and conclusions.

II. GRAPH SIGNALS, CONVOLUTIONAL OPERATORS ON GRAPHS, AND GNNS

In this section we provide a basic description of signals and convolutional operators on graphs. Additionally, we discuss GNNs and their mapping operators.

A. Graph signal processing

Let us consider the graph $G = (V(G), E(G), w_G)$ with set of vertices $V(G)$, set of edges $E(G)$, and weight function
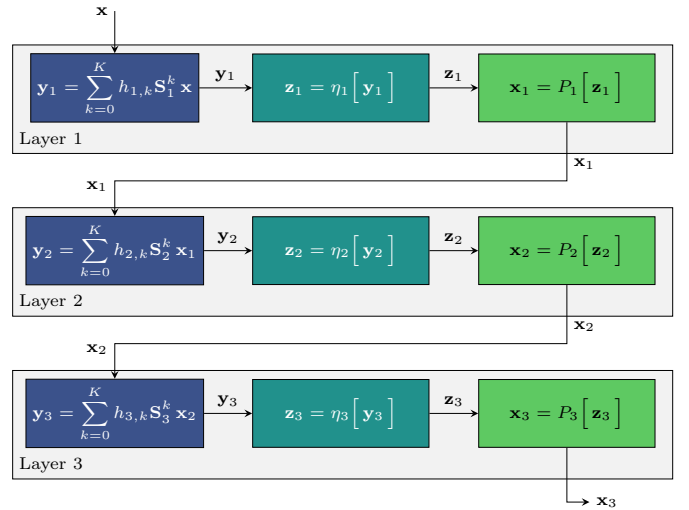

Figure 1. Graph neural network with three layers. The input signal $\mathbf{x}$ is processed by the GNN to produce an output $\mathbf{x}_3$. In the i-th layer the information is processed by a convolution operator $\sum_{k=0}^{K} h_{i,k} \mathbf{S}_i^k$, then by a pointwise nonlinearity η_i, and then by a pooling operator P_i that matches the dimensions of the signals between layers.

$w_G : E(G) \rightarrow \mathbb{R}^+$. We define a graph signal $\mathbf{x}$ on G as the map $\mathbf{x} : V(G) \rightarrow \mathbb{R}$ identified with a vector in $\mathbb{R}^{|V(G)|}$. The i-th component of $\mathbf{x}$ is the value of $\mathbf{x}$ on the i-th node in $V(G)$, given an ordering of $V(G)$. In what follows we will use the symbol $(G, \mathbf{x})$ to denote the graph signal $\mathbf{x}$ defined on G. We recall that the image set of $E(G)$ under w_G can be stored in a weight or adjacency matrix $\mathbf{A}$, with $\mathbf{A}(i,j) = w_G(\{i,j\})$, $\{i,j\} \in E(G)$.

In graphs, the notion of filtering and convolution with filters relies on a shift operator $\mathbf{S} \in \mathbb{R}^{N \times N}$ with $N = |V(G)|$, which can be selected as the Laplacian matrix, normalized Laplacian, or the adjacency matrix $\mathbf{A}$ [19], [21], [22]. In this paper we consider $\mathbf{S} = \mathbf{A}$ as the shift operator. Then, filters on the graph are polynomial matrix operators $\mathbf{H}(\mathbf{S}) = \sum_{k=0}^{K-1} h_k \mathbf{S}^k$, where the coefficients h_k are called the filter taps. The convolution between a filter $\mathbf{H}(\mathbf{S})$ and a signal $(G, \mathbf{x})$ results in a signal $(G, \mathbf{y})$ with

$$\mathbf{y} = \mathbf{H}(\mathbf{S})\mathbf{x} = \sum_{k=0}^{K-1} h_k \mathbf{S}^k \mathbf{x}. \tag{1}$$

B. Graph neural networks

A graph neural network (GNN) is a stacked layered structure – see Fig. 1. In each layer, information is processed by means of convolutional operators acting on graph signals, followed by pointwise nonlinearities and pooling operators. Then, in the i-th layer of the GNN an input signal $\mathbf{x}_{i-1}$ is filtered as $\mathbf{y}_i = \sum_{k=0}^{K} h_{i,k} \mathbf{S}_i^k \mathbf{x}_{i-1}$. Afterwards, a point-wise nonlinearity operator η_i is applied to $\mathbf{y}_i$ to obtain $\mathbf{z}_i = \eta_i(\mathbf{y}_i)$. Finally, a pooling operator P_i reduces the dimension of $\mathbf{z}_i$ and the size of G_i – which implies changes in $\mathbf{S}_i$ – generating $\mathbf{x}_i = P_i(\mathbf{z}_i)$. The operation defined by P_i is meant to preserve structural properties of information, and the function η_i is required to be Lipschitz [5], [23], [29]. The output of the i-th layer can

be written as $\mathbf{x}_i = P_i\left(\eta_i\left(\mathbf{H}_i(\mathbf{S}_i)\mathbf{x}_{i-1}\right)\right)$, with $\mathbf{H}_i(\mathbf{S}_i) = \sum_{k=0}^{K} h_{i,k}\mathbf{S}_i^k$. In each layer several filters can be considered to produce outputs of multiple features.

Since the pooling operator P_i reduces the dimensionality of the data, this entails a possible modification of the underlying graphs in the layers. Among the pooling methods considered for GNNs two approaches stand out, graph coarsening and zero padding. When graph coarsening is used, spectral clustering techniques are used to group subsets of nodes and edges in a graph G_i to define a new graph G_{i+1} where $|V(G_{i+1})| < |V(G_i)|$ [24], [25]. To do graph pooling with zero padding, the dimensions of the graphs are preserved while some components of the signals are forced to be zero [23]. To optimally select those components, we use graph sampling approaches like those proposed in [8]–[17]. Then, while zero padding does not modify the dimensions of the graphs, i.e. $|V(G_{\ell+1})| = |V(G_\ell)|$, it does implicitly modify the shift operators, which are now associated to an *induced subgraph* of the original graph.

The coefficients of the convolutional operators h_k are learned from the data. This is, for a training set $\mathcal{T} = \{(\mathbf{x}, \mathbf{y})\}$ with inputs $\mathbf{x}$ and outputs $\mathbf{y}$, the GNN learns a representation mapping that associates an output $\hat{\mathbf{y}}$ for a given input $\hat{\mathbf{x}}$ with $(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \notin \mathcal{T}$ [23].

To represent the mapping operator of a GNN with L layers we use the symbol $\boldsymbol{\Phi}\left(\mathbf{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\mathbf{S}_\ell\}_{\ell=1}^L\right)$, where $\mathbf{x}$ is the input signal to the GNN, $\mathcal{P}_\ell$ is the set of filters used in the ℓ-th layer, and $\mathbf{S}_\ell$ is the ℓ-th shift operator.

III. SIGNALS AND CONVOLUTIONAL OPERATORS ON GRAPHONS

A graphon is a symmetric bounded function $W : [0,1]^2 \to [0,1]$ that can be conceived as a *completion* of the limits for infinite sequences of graphs in the same way that irrational numbers complete the real line [26], [30], [31]. Some graphons can be obtained from labeled graphs. For instance, given a graph G with weight matrix $\mathbf{A}$ one can obtain an *induced graphon* W_G given by $W_G(x,y) = \mathbf{A}(\lceil Nx \rceil, \lceil Ny \rceil)$, where $\lceil \cdot \rceil$ is the ceiling operator, $N = |V(G)|$, and $x, y \in [0,1]$. Therefore, W_G is a piece-wise constant function on a regular grid, and the amplitudes of W_G are given by the weight matrix $\mathbf{A}$.

A. Convergence of graphs to graphons

The concept of convergence of sequences of graphs to graphons relies on the convergence of sequences of homomorphism densities – see subsection A-A3 in the Appendix. This type of convergence is compatible with a notion of convergence based on the cut norm $\|\cdot\|_\square$. To see this, let us consider the metric distance $\delta_\square(W_{G_i}, W)$ between W_{G_i} and W given by $\delta_\square(W_{G_i}, W) = \inf_\pi \|W - \pi(W_{G_i})\|_\square$, where $\pi : [0,1] \to [0,1]$ is a measure preserving map – analogous to a node label permutation on the graph –, and the cut norm is given by

$$\|W(x,y)\|_\square = \sup_{I,J \subset [0,1]} \left| \int_{I \times J} W(x,y)dxdy \right|. \quad (2)$$

As shown in [26], [30], if $\{G_i\}_i$ converges to $W(x,y)$ in the homomorphism sense – which we denote by $\{G_i\} \to W(x,y)$ –, then $\delta_\square(W_{G_i}, W) \to 0$. Notice that the action of π on $W_{G_i}(x,y)$ is given by $\pi(W_{G_i}(x,y)) = W_{G_i}(\pi(x), \pi(y))$.

B. Convolutional information processing on graphons

The connection between graphs and graphons prompts a natural link between convolutional signal processing on graphs and convolutional signal processing on graphons. We begin by recalling that in graphon signal processing (Gphon-SP) a signal on a graphon $W(x,y)$ is given by a pair $(W, \boldsymbol{x})$ where $\boldsymbol{x}$ is an element of $L^2([0,1])$.

In the same way that graphons can be induced by graphs, some graphon signals can be induced by graph signals. To see this, let us consider the graph signal $(G, \mathbf{x})$ and the graphon signal $(W_G, \boldsymbol{x})$ where W_G is induced by G. Then, we can say that $(W_G, \boldsymbol{x})$ is induced by $(G, \mathbf{x})$ if $\boldsymbol{x}(t) = \mathsf{step}(\mathbf{x}) = \mathbf{x}(\lceil tN \rceil)$, where $t \in [0,1]$. This is, graphon signals induced by graph signals are piece-wise constant functions defined on a regular partition of $[0,1]$. Notice that a graphon signal is a particular case of a *node level statistic* as defined in [26].

To define convolutions for graphon signals, it is necessary to first introduce a shift operator [20], [32], [33]. Following [30], we define the graphon shift operator $\boldsymbol{T}_W$ on a graphon $W(x,y)$ acting on a graphon signal $(W, \boldsymbol{x})$ by

$$(\boldsymbol{T}_W \boldsymbol{x})(v) = \int_0^1 W(u,v)\boldsymbol{x}(u)du. \quad (3)$$

Notice that $\boldsymbol{T}_W$ is Hilbert-Schmidt [30]. Now, we can define convolutional filters for graphon signals using polynomial operators written in terms of $\boldsymbol{T}_W$. The convolutional filtering of a graphon signal $(W, \boldsymbol{x})$ by means of a graphon filter $h(\boldsymbol{T}_W) = \sum_{k=0}^K h_k \boldsymbol{T}_W^k$ is given by

$$h(\boldsymbol{T}_W)\boldsymbol{x} = \sum_{k=0}^K h_k \boldsymbol{T}_W^k \boldsymbol{x}, \quad (4)$$

where $\boldsymbol{T}_W^k$ represents the k-times composition of $\boldsymbol{T}_W$. It is important to remark that the graphon filter $h(\boldsymbol{T}_W)$ can be characterized by means of the scalar polynomial [32], [33],

$$h(t) = \sum_{k=0}^K h_k t^k. \quad (5)$$

In what follows we will refer to $h(t)$ as the polynomial or functional representation of the graphon filter $h(\boldsymbol{T}_W) = \sum_{k=0}^K h_k \boldsymbol{T}_W^k$. The representation $h(t)$ is also known as the frequency representation of the graphon filter [32], [33].

The relationship between graph signals and their induced graphon counterparts translate to the action of the shift operators and the filters. However, as we will indicate in the theorem below, that transference requires a scaling.

Theorem 1. *Let $(W_G, \boldsymbol{x})$ be a graphon signal induced by the graph signal $(G, \mathbf{x})$. Let $h(t) = \sum_{k=0}^{K-1} h_k t^k$ be a filter and $\boldsymbol{y} = h(\boldsymbol{T}_{W_G})\boldsymbol{x}$, where $\boldsymbol{T}_{W_G}$ is the graphon shift operator in W_G. Then, it follows that*

$$\boldsymbol{y} = \mathsf{step}\left(h\left(\frac{\mathbf{S}_G}{|V(G)|}\right)\mathbf{x}\right), \quad (6)$$

where $\mathbf{S}_G$ *is the shift operator on* G.

Proof. See Appendix B-A ☐

From Theorem 1 we observe that in order to transfer the filtering operation between graph signals and their induced graphon counterparts, we have to scale the shift operator in the graph according to $\mathbf{S}_G \rightarrow \mathbf{S}_G/|V(G)|$. Notice that although $\boldsymbol{y} \neq \text{step}(h(\mathbf{S}_G)\mathbf{x})$, it is possible to determine uniquely $\boldsymbol{y}$ from $\text{step}(h(\mathbf{S}_G)\mathbf{x})$ given the knowledge of $h(t)$. This implies that it is possible to learn the coefficients of a filter on the graph and transfer them to the induced graphon. Likewise, it is possible to characterize the properties of the filters on the induced graphon operators and transfer such properties to the filters on the graph. We will exploit this relationship in order to analyze GNNs. More specifically, we characterize the properties of the graph and GNN operators on their induced graphon representations.

C. Graphon Neural Networks (Gphon-NNs)

Convolutional architectures can be established in multiple domains including graphons [32]. This is, information processing on graphons can be combined with pointwise nonlinearity operators to obtain *graphon neural networks (Gphon-NNs)* [1]. Formally, a Gphon-NN is a stacked layered structure analogous to a GNN – see Fig. 1 – where the convolutions are carried out by graphon filters on graphon signals. In the i-th layer of a Gphon-NN an input signal $\boldsymbol{x}_{i-1}$ is filtered to obtain $\boldsymbol{y}_i = \sum_{k=0}^K h_{i,k}\boldsymbol{T}_{W_i}^k\boldsymbol{x}_{i-1}$. Then, $\boldsymbol{y}_i$ is transformed by a pointwise nonlinearity η_i – assumed to be Lipschitz – to obtain $\boldsymbol{z}_i = \eta_i(\boldsymbol{y}_i)$ –, and right after a pooling operator generates $\boldsymbol{x}_i = P_i(\boldsymbol{z}_i)$. The operator P_i reduces the complexity and or degrees of freedom of the information in $\boldsymbol{z}_i$ while preserving essential features. Additionally, P_i can be embedded in the graphon shift operator associated to each layer. Then, the output of the i-th layer of a Gphon-NN can be written as $\boldsymbol{x}_i = \eta_i(\boldsymbol{H}_i(\boldsymbol{T}_{W_i})\boldsymbol{x}_{i-1})$, with $\boldsymbol{H}_i(\boldsymbol{T}_{W_i}) = \sum_{k=0}^K h_{i,k}\boldsymbol{T}_{W_i}^k$. This expression can be extended trivially to multiple features but such extension is not central to our analysis.

Similar to the scenario of GNNs, the coefficients $h_{i,k}$ are learned from the input data. Given a training set $\mathcal{T} = \{(\boldsymbol{x}, \boldsymbol{y})\}$ with inputs $\boldsymbol{x}$ and outputs $\boldsymbol{y}$, the Gphon-NN learns a representation that relates an output $\hat{\boldsymbol{y}}$ to a given input $\hat{\boldsymbol{x}}$ with $(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}) \notin \mathcal{T}$. We will represent the mapping operator of a Gphon-NN with L layers by $\boldsymbol{\Phi}\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\boldsymbol{T}_{W_\ell}\}_{\ell=1}^L\right)$. The input signal to the Gphon-NN is $\boldsymbol{x}$, while $\mathcal{P}_\ell$ and $\boldsymbol{T}_{W_\ell}$ are the sets that indicate the properties of the filters and the graphon shift operators in the layer ℓ, respectively.

Remark 1. Notice that as a consequence of the unique relationship between graphs and their induced graphons, it is possible to use Gphon-NNs to process information on GNNs. To see this, let us consider the GNN with graph layers given by $\{G_\ell\}_{\ell=1}^L$. Then, there is an induced Gphon-NN with graphon layers given by $\{W_{G_\ell}\}_{\ell=1}^L$. Let $(W_{G_\ell}, \boldsymbol{x})$ be the graphon signal induced from $(G_\ell, \mathbf{x})$. Then, the processing of $(G_\ell, \mathbf{x})$ in the ℓ-th layer of the GNN can be carried out by processing the graphon signal $(W_{G_\ell}, \boldsymbol{x})$ in the ℓ-th layer of the induced Gphon-NN – taking into account the scaling given in
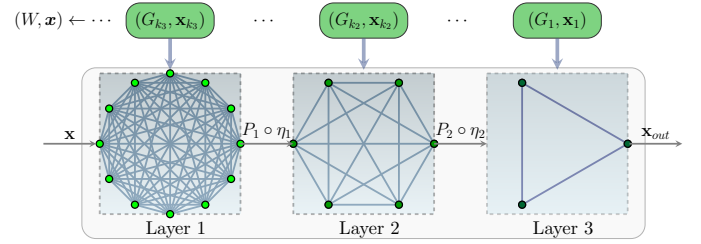

Figure 2. Schematic representation of the graphon pooling concept. The sequence of graph signals $(G_i, \mathbf{x}_i)$ converge to the graphon signal $(W, \boldsymbol{x})$. Then, a finite subsequence of graph signals $\{(G_1, \mathbf{x}_1), (G_{k_2}, \mathbf{x}_{k_2}), (G_{k_3}, \mathbf{x}_{k_3})\}$ is selected to define the layers of the GNN, where the largest graph is associated to the first layer. This process specifies in an implicit way the action of the pooling operators P_i between layers since the convergence of $(G_i, \mathbf{x}_i)$ guarantees some minimum common structural properties for all the elements in the sequence.

Theorem 1. Then, there is a one to one relationship between $\boldsymbol{\Phi}\left(\mathbf{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\mathbf{S}_\ell\}_{\ell=1}^L\right)$ and $\boldsymbol{\Phi}\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\boldsymbol{T}_{W_{G_\ell}}\}_{\ell=1}^L\right)$. Additionally, if the graphs $\{G_\ell\}_{\ell=1}^L$ are obtained from $W(x, y)$ by a pooling method, the effects of pooling on the GNN can be studied analyzing the term $\left\|\boldsymbol{\Phi}\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \boldsymbol{T}_W\right) - \boldsymbol{\Phi}\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\boldsymbol{T}_{W_{G_\ell}}\}_{\ell=1}^L\right)\right\|_2$. This means that the effects of the pooling operation can be analyzed by means of a comparison between the Gphon-NN induced from the GNN, and a Gphon-NN defined on the graphon $W(x, y) - W(x, y)$ is associated to the graph in the first layer of the GNN. In the following section we leverage these facts to study the effects of the graphon pooling methods proposed.

IV. GRAPHON POOLING

The fundamental idea of the graphon pooling approach proposed in this paper relies on leveraging subsequences of a convergent sequence of graph signals to build a GNN – see Fig. 2 . The largest graph of such subsequence is associated to the first layer of the GNN, while the smaller graphs are associated to subsequent layers. The convergence of the graph signals implies spectral consistency on the graphs – that are also convergent – and that there is functional convergence of the information associated to the nodes of the graphs. The sequences are built from the data embedded in a graphon representation.

It is important to remark that in practice one starts with the data defined on a large graph and not a graphon. However, as we have shown that every graph has a natural induced graphon representation, we will show in subsection IV-B that such induced graphon representation is a good approximation of the graphon limit if the graph is large enough. This indicates that the original graph provides naturally the graphon used to build the elements of the convergent graph sequence. In what follows we describe the numerical approaches used to build such sequences.

Method 1 (M1, regular integration). In this approach we start by considering a graphon $W(u, v)$. We then build a sequence of graphs $\{G_\ell\}_\ell \rightarrow W(u, v)$ using a uniform partition of

$[0,1]^2 \subset \mathbb{R}^2$ – see Fig. 3 (left). We associate the entries of the adjacency matrix of the graph with the volume below $W(u,v)$ in each of the elements of the partition. Then, the adjacency matrix $\mathbf{A}^{G_\ell}$ for each graph G_ℓ is given by

$$\mathbf{A}^{G_\ell}(i,j) = \frac{1}{\Delta_{i,j}} \int_{\rho(j)}^{\rho(j+1)} \int_{\rho(i)}^{\rho(i+1)} W(u,v)dudv, \quad (7)$$

with $\rho(i) = (i-1)/N_\ell$ for all $i \in \{1,2,\dots N_\ell\}$, $\Delta_{i,j} = 1/N_\ell^2$, where N_ℓ is the number of nodes in G_ℓ, and $\rho(N_\ell + 1) = 1$. The schematic representation of this method is depicted in Fig. 3 (left).

Method 2 (**M2**, irregular integration). In this approach we build the graphs in the sequence considering an irregular partition of $[0,1]^2 \subset \mathbb{R}^2$ – see Fig. 3 (center). We associate each entry of the adjacency matrix to the volume below $W(x,y)$ in each element of the partition. To build the partition we use a random uniform distribution over $[0,1]$. The adjacency matrix $\mathbf{A}^{G_\ell}$ of each G_ℓ is given by (7), where $\Delta_{i,j}$ is the area of the partition element $[\rho(j), \rho(j+1)] \times [\rho(i), \rho(i+1)]$, and the map ρ is defined according to a random uniform distribution in $[0,1]$ with the following restrictions: $\rho(1) = 0$, $\rho(i) \leq \rho(i+1)$, and $\rho(N_\ell + 1) = 1$. This procedure is illustrated in Fig. 3 (center).

Method 3 (**M3**, irregular sampling). In this approach we obtain the adjacency matrix of the graphs in the sequence sampling values of an underlying graphon $W(x,y)$. To perform the sampling we use a random uniform distribution to select a set of N_ℓ points in $[0,1]$. The relationship between the nodes and the points in $[0,1]$ is defined by a map ρ with $\rho(i) \leq \rho(i+1)$. Then, the adjacency matrix $\mathbf{A}^{G_\ell}$ of the graph G_ℓ is given by

$$\mathbf{A}^{G_\ell}(i,j) = W(\rho(i), \rho(j)). \quad (8)$$

The discretization method **M3** is illustrated in Fig. 3 (right).

A. Labeling and mapping of signal components

The labeling assigned to the graph $\mathbf{A}^G$ generated by means of **M1**, **M2**, and **M3** is inherited from the natural ordering associated to the interval $[0,1]$. To see this we recall that for all the discretization methods proposed we have $\rho(i) < \rho(i+1)$. Then, the index i naturally *labels* the nodes in $\mathbf{A}^G$ and the components of any graph signal defined on $\mathbf{A}^G$.

Methods **M1**, **M2**, and **M3** do not establish a correspondence between components of signals on graphs of different sizes. Since **M1/M2/M3** are used to define the layers of a GNN, a consistent mapping between the signal components is necessary. We perform such mapping leveraging the associated induced graphon signals and interpolation. The details of this procedure are presented in the following paragraphs.

Signal mapping with M1/M2: In order to perform the interpolation of signals when **M1/M2** is used, we take as a reference the center point of the intervals defining the partition used to build the integration grid. Let us consider the pooling operation on the graph G_ℓ to obtain $G_{\ell+1}$. We denote by $I_i^{(\ell)}$
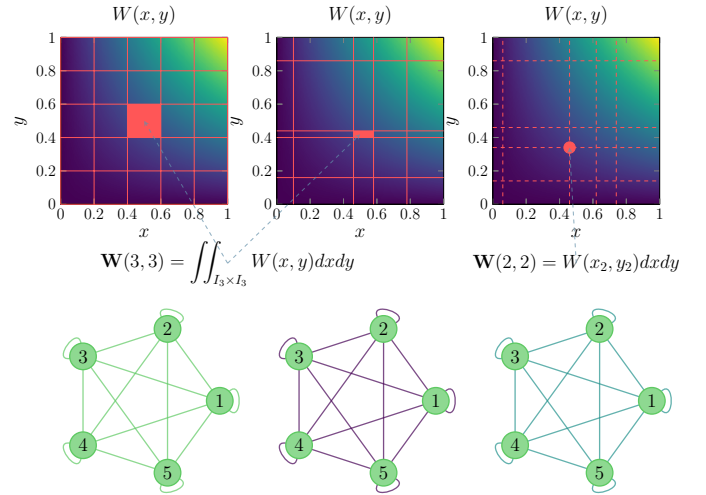

Figure 3. Generation of a graph adjacency matrix $\mathbf{A}^G$ from an underlying Graphon $W(x,y)$ by means of methods **M1** (left), **M2** (center), and **M3** (right). The grids in the first row (left and center) define partitions of $[0,1]^2$, while the dotted lines in the right (first row) indicate values of x and y where $W(x,y)$ is evaluated. Left: discretization method **M1** where a *regular* grid is used to build $\mathbf{A}^G$, and each entry of $\mathbf{A}^G$ is given by the volume below $W(x,y)$ in each element of the partition. Center: discretization method **M2**, where the subdomains of an *irregular* grid are used to obtain the entries of $\mathbf{A}^G$. Right: discretization method **M3** where each entry of $\mathbf{A}^G$ is obtained evaluating $W(x,y)$ at an specific point $(x,y) \in [0,1]^2$.

the intervals of the partition of $[0,1]$ associated to W_{G_ℓ} and by $\overline{I_i^{(\ell)}}$ the center point of $I_i^{(\ell)}$. Now, we take into account that for any $I_j^{(\ell+1)}$ there exists $I_i^{(\ell)}$ and $I_{i+1}^{(\ell)}$ such that

$$\overline{I_i^{(\ell)}} < \overline{I_j^{(\ell+1)}} < \overline{I_{i+1}^{(\ell)}}. \quad (9)$$

Then, if the graphs G_ℓ are obtained from $W(x,y)$ using **M1/M2** the signal $(W_{G_\ell}, \boldsymbol{x}_\ell)$ is mapped to the signal $(W_{G_{\ell+1}}, \boldsymbol{x}_{\ell+1})$, where

$$\boldsymbol{x}_{\ell+1}(u) = \frac{1}{2}\left(\boldsymbol{x}_\ell\left(I_i^{(\ell)}\right) + \boldsymbol{x}_\ell\left(I_{i+1}^{(\ell)}\right)\right), \quad u \in I_j^{(\ell+1)}. \quad (10)$$

Since $\boldsymbol{x}_\ell$ is piece-wise constant on the intervals $I_i^{(\ell)}$, the value $\boldsymbol{x}_\ell\left(I_i^{(\ell)}\right)$ is well defined.

Signal mapping with M3: For the interpolation using **M3** we use the location of the sampling points. We perform the pooling operation on the graph G_ℓ to obtain $G_{\ell+1}$ and we denote by $t_i^{(\ell)}$ the sampling points in $[0,1]$ associated to W_{G_ℓ}. We recall that for any $t_j^{(\ell+1)}$ there exists $t_i^{(\ell)}$ and $t_{i+1}^{(\ell)}$ such that

$$t_i^{(\ell)} < t_j^{(\ell+1)} < t_{i+1}^{(\ell)}. \quad (11)$$

Therefore, if the graphs $\{G_\ell\}_\ell$ are obtained from $W(x,y)$ by means of method **M3** the signal $(W_{G_\ell}, \boldsymbol{x}_\ell)$ is mapped to the signal $(W_{G_{\ell+1}}, \boldsymbol{x}_{\ell+1})$, where

$$\boldsymbol{x}_{\ell+1}\left(t_j^{(\ell+1)}\right) = \frac{1}{2}\left(\boldsymbol{x}_\ell\left(t_i^{(\ell)}\right) + \boldsymbol{x}_\ell\left(t_{i+1}^{(\ell)}\right)\right). \quad (12)$$

B. Consistency of graphon pooling

In this section we elaborate about the theoretical foundations of graphon pooling. We start with introducing a result providing the convergence guarantees for those sequences of graphs obtained from a graphon using the discretization methods introduced before. In what follows we will refer to discretization methods **M1**, **M2** and **M3** as the pooling methods.

Theorem 2. *Let $\{G_\ell\}_\ell$ be a sequence of graphs generated from the graphon $W(x, y)$ by the discretization methods **M1** and **M2**. If $\{W_{G_\ell}\}_\ell$ is the sequence of induced graphons associated to $\{G_\ell\}_\ell$, then $\{W_{G_\ell}\}_\ell \to W$ almost everywhere with*

$$\left\| \boldsymbol{T}_W - \boldsymbol{T}_{W_{G_\ell}} \right\|_{\infty \to 1} \leq \frac{8}{\sqrt{\log |V(G_\ell)|}}. \quad (13)$$

Proof. See Appendix B-B. $\qquad\square$

Theorem 2 provides the guarantees of convergence of a sequence of graphs used to build a large GNN. This result at the same time assures the spectral structural consistency of the graphs used in the GNN. We emphasize that Theorem 2 is an immediate consequence of the fact that kernels can be approximated well by step functions in the L_1 norm – see Appendix B-B.

In graphon pooling the assumption that there is a closed form expression for the limit graphon can indeed be relaxed when considering a piecewise constant representation of the largest finite graph in the sequence. This is not just an approximation, but as we will show in the following lemma it is indeed a natural way to represent a graphon with zero error in terms of the cut metric.

Lemma 1 (Adapted from Lemma 9.11 in [30])**.** *Let G be a graph obtained from the graphon $W(x, y)$ using the discretization methods **M1** and **M2**. Then, there exists a graph H with $|V(H)| \leq 4|V(G)|$ such that*

$$\|W - W_G\|_\square = \|W_H - W_G\|_\square, \quad (14)$$

*where H is obtained from $W(x, y)$ using methods **M1** and **M2** with a refined partition of W_G. The terms W_G and W_H are the graphons induced by G and H, respectively.*

Lemma 1 highlights fundamental properties of the graphon pooling methods proposed in this paper. In particular, it remarks that in terms of the cut norm a step function representation of any graphon is enough to capture structural properties. In practice this points to the fact that a large graph induces a graphon that is a good approximation of the graphon limit. Then, if a finite sequence of graphs in a GNN is obtained from a graphon, one could consider that the largest graph in the sequence can induce a graphon that is indistinguishable from the graphon limit – in terms of the cut norm. It is equally important to highlight that the partition associated to W_H has at most $4|V(G)|$ elements. In practice this implies that one can reduce the size of a step function graphon up to a factor of 4 and still preserve the same distance with respect to the graphon limit.

Taking into account the cut norm distance between the induced graphons of a convergent sequence $\{W_{G_\ell}\}_\ell$ and the graphon limit $W(x, y)$, we can obtain upper bounds for the change of the filter outputs when implemented on W_{G_ℓ} and W. We state this formally in the following theorem.

Theorem 3. *Let $h(t) = \sum_{k=0}^\infty h_k t^k$ be the functional representation of a graph/graphon filter with Lipschitz constant L. Let G be a graph obtained from the graphon $W(x, y)$ using methods **M1** or **M2**. Let $\|\boldsymbol{T}_W - \boldsymbol{T}_{W_G}\|_{HS} \leq \gamma \|\boldsymbol{T}_W - \boldsymbol{T}_{W_G}\|_2$ for $\gamma > 0$, where $\boldsymbol{T}_W$ and $\boldsymbol{T}_{W_G}$ are the graphon shift operators of W and W_G respectively. Then, there exists a graph H obtained from $W(x, y)$ using **M1** or **M2** such that*

$$\|h(\boldsymbol{T}_W) - h(\boldsymbol{T}_{W_G})\|_2 \leq$$
$$\gamma L \sqrt{8 \|W_H - W_G\|_\square} + \mathcal{O}\left(\|\boldsymbol{T}_W - \boldsymbol{T}_{W_G}\|_2^2\right), \quad (15)$$

with $|V(H)| \leq 4|V(G)|$, and where $\|\cdot\|_2$ represents the norm $\|\cdot\|_{2 \to 2}$ and $\|\cdot\|_{HS}$ indicates the Hilbert-Schmidth norm.

Proof. See Appendix B-D. $\qquad\square$

From Theorem 3 we can see that the size of the difference between the outputs of filters instantiated on $W(x, y)$ and W_G has an upper bound that can be calculated by means of graphs that can be obtained from $W(x, y)$ using methods **M1** or **M2**. This is, the square root of the distance between $W(x, y)$ and W_G in terms of the cut norm completely determines the distance between convolutional filters on both graphons in the operator norm.

With this result at hand, we now establish an upper bound for the change of the GNN and Gphon-NN operators when considering graphs obtained from a given graphon using methods **M1** or **M2**.

Theorem 4. *Let $\Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \boldsymbol{T}_W\right)$ be the mapping operator of a Gphon-NN where no pooling operation is performed. Let $\Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\boldsymbol{T}_{W_{G_\ell}}\}_{\ell=1}^L\right)$ be the mapping operator of a Gphon-NN where each layer is defined by the graphon W_{G_ℓ}, where the G_ℓ are generated from $W(x, y)$ by methods **M1/M2**. Let $\|\boldsymbol{T}_W - \boldsymbol{T}_{W_{G_\ell}}\|_{HS} \leq \gamma \|\boldsymbol{T}_W - \boldsymbol{T}_{W_{G_\ell}}\|_2$ for $\gamma > 0$, where $\boldsymbol{T}_W$ and $\boldsymbol{T}_{W_{G_\ell}}$ are the graphon shift operators of W and W_{G_ℓ}, respectively. Then, it follows that*

$$\left\| \Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \boldsymbol{T}_W\right) - \Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\boldsymbol{T}_{W_{G_\ell}}\}_{\ell=1}^L\right) \right\|_2$$
$$\leq \sqrt{8} L \gamma \left(\sum_{\ell=1}^L \sqrt{\|W_{H_\ell} - W_{G_\ell}\|_\square} \right) \|\boldsymbol{x}\|_2$$
$$+ \mathcal{O}\left(\|\boldsymbol{T}_W - \boldsymbol{T}_{W_{G_\ell}}\|_2^2\right), \quad (16)$$

*where $\mathcal{P}_\ell$ is a set of L-Lipschitz filters, H_ℓ is obtained from $W(x, y)$ by **M1/M2** and $|V(H_\ell)| \leq 4|V(G_\ell)|$. The index (ℓ) makes reference to quantities and constants associated to the layer ℓ.*

Proof. See Appendix B-E. $\qquad\square$

Theorem 4 highlights that the pointwise nonlinearities η_ℓ in the Gphon-NNs do not contribute to increase the difference between $\Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \boldsymbol{T}_W\right)$ and

$\Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\boldsymbol{T}_{W_{G_\ell}}\}_{\ell=1}^L\right)$. This means that, any difference between an induced Gphon-NN from a GNN and its ideal Gphon-NN version on $W(x,y)$ is due to the operation of pooling, and this is captured in the upper bound in (16) stated in terms of the cut norm between two induced graphons. Additionally, the value of this bound can be calculated with graphs obtained from $W(x,y)$ using methods **M1** and **M2**.

Remark 2. Notice that the condition of being Lipschitz for the filters in Theorems 3 and 4 is analog to the condition of being low pass for filters in discrete signal processing (DSP). While the high frequencies in DSP are associated to the smallest eigenvalues, high frequencies for graphon signals are associated to the lowest eigenvalues since cero is an accumulation point of the graphon shift operator. In fact, due to the nature of the spectrum of the graphon shift operator if the graphon filter $h(t)$ is Lipschitz it must be flatten out towards $t=0$.

Now we turn our attention to error estimates when we start with a graph approximation of an underlying graphon – obtained by unspecified means – and we perform the operation of pooling on those graph estimates.

Theorem 5. *Let $W(x,y)$ be a graphon and let G_1 and G_2 be two graph estimates of W such that $\|\boldsymbol{T}_W - \boldsymbol{T}_{W_{G_i}}\|_{HS} \leq \epsilon$. Let H_i be the graph obtained from W_{G_i} by method **M1**. If $\|\boldsymbol{T}_{W_{G_i}} - \boldsymbol{T}_{W_{H_i}}\|_{\infty\to 1} \leq \epsilon/V(H_i)^4$ we have that*

$$\left\|\boldsymbol{T}_{W_{H_1}} - \boldsymbol{T}_{W_{H_2}^\theta}\right\|_{\infty\to 1} \leq 32\epsilon, \tag{17}$$

where $W_{H_2}^\theta = W_{H_2}(\theta(x), \theta(y))$ and θ is any measure preserving map on $[0,1]$.

Proof. See Appendix B-F. $\qquad\square$

Notice that the graphs G_1 and G_2 are discrete approximations of $W(x,y)$, obtained by arbitrary means. Then, Theorem 5 shows that when the estimates G_1 and G_2 of $W(x,y)$ lead to $\boldsymbol{T}_W$ and $\boldsymbol{T}_{W_{G_i}}$ that are ϵ-close in the Hilbert-Schmidt norm, we can apply the pooling method **M1** on W_{G_i} to obtain graphs H_i such that $\boldsymbol{T}_{W_{H_1}}$ and $\boldsymbol{T}_{W_{H_2}}$ are also ϵ-close. This indicates that if two estimates of a graphon are close, the pooling method **M1** preserves this closeness. This fact endows the pooling method **M1** with a stability that is not guaranteed for the pooling approaches **M2** and **M3**, which also provides the theoretical support for its better performance shown in numerical simulations.

Remark 3. In the following section we perform a set of numerical experiments where we evaluate graphon pooling in GNNs. We remark that although the results derived above are expressed in terms of Gphon-NNs, they are directly associated to those GNNs that induce the Gphon-NNs. As we have stressed before in Remark 1 there is a one to one relationship between the GNN with mapping operator $\Phi\left(\mathbf{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\mathbf{S}_\ell\}_{\ell=1}^L\right)$ and its induced Gphon-NN with mapping operator given by $\Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\boldsymbol{T}_{W_{G_\ell}}\}_{\ell=1}^L\right)$. Then, we evaluate the proposed graphon pooling methods with GNNs where graphs are constructed with the discretization

methods (**M1**, **M2** and **M3**). Together with the numerical validations in the following section, the effectiveness of the graphon pooling operation can be proved both theoretically and numerically.

V. NUMERICAL SIMULATIONS

We consider three problem settings to verify the performance of our proposed graphon pooling methods, i.e. **M1**(regular integration), **M2**(irregular integration) and **M3**(irregular sampling). We also consider a regular sampling method which is similar to **M3** but with a uniform grid as sampling points. By comparing with graph coarsening [24] and selection GNNs with zero padding [23], we claim that graphon pooling can achieve better performance while spending less time. All of the architectures are trained in parallel implementing the ADAM algorithm for stochastic optimization [34] with decaying factors set as $\beta_1 = 0.9$ and $\beta_2 = 0.999$.

A. Source localization problem

We synthetically generate diffusion processes on graphs acquired from several graphon functions. The graph has N nodes with the shift operator $\mathbf{S} \in \mathbb{R}^{N \times N}$ and let $\mathbf{x}_0 \in \mathbb{R}^N$ be a graph signal such that $[\mathbf{x}_0]_i = 1$ for node $i = c$ and 0 otherwise. The graph diffusion process $\{\mathbf{x}_t\}$ is defined as

$$\mathbf{x}_t = \mathbf{S}^t \mathbf{x}_0 \tag{18}$$

which represents the t-step diffusion graph signal. The objective is to locate the source node c given $\mathbf{x}_t$ for arbitrary t. The source node c is selected among $C = 10$ possible sources, which makes this source localization problem as a classification problem on $C = 10$ classes.

We train all the architectures with different pooling strategies to solve this problem by minimizing the cross-entropy loss on $1,000$ $(\mathbf{x}_t, c)$ training samples with learning rate 5×10^{-4}. The training dataset is divided in batches of 20, over 300 epochs. The learned architectures are validated and tested by evaluating the classification error rates on sets with 240 and 200 samples respectively. All the architectures are made up of GNNs with $L = 2$ layers. Each layer contains $F_1 = F_2 = 8$ output features and $K_1 = K_2 = 5$ filter taps.

We first analyze the performances of different graphon pooling methods on graphs obtained from 3 graphon models: the exponential graphon $W(x,y) = \exp(-\beta(x-y)^2)$ with $\beta = 2.3$; the bilinear graphon $W(x,y) = xy$; and the polynomial graphon $W(x,y) = 0.5(x^2 + y^2)$. The initial graphs with $N = 100$ nodes are generated from function $W(x,y)$ by integration or sampling either regularly or irregularly. The number of selected nodes in the first and second layers in all architectures are $N_1 = 50$ and $N_2 = 25$ respectively. The layer-wise dimensionality reduction ratios are fixed as 2 because of the setting of graph coarsening algorithm. The test classification error rates achieved with different graphon pooling methods, graph coarsening and the selection GNN are presented in Table I for exponential, bilinear and polynomial graphons. We report the average error rate and standard deviation for models trained on 8 different dataset realizations.

Architecture	$W(x,y)$		
	$e^{-\beta(x-y)^2}$	xy	$0.5(x^2+y^2)$
M1 pooling	10.88 ± 2.52	7.00 ± 4.47	7.25 ± 5.55
Graph coarsening	15.69 ± 4.78	6.56 ± 3.55	7.10 ± 4.39
Selection GNN	21.69 ± 6.95	26.12 ± 12.61	8.05 ± 5.58

(a) Source localization with graphs initiated with **M1** graphon pooling

Architecture	$W(x,y)$		
	$e^{-\beta(x-y)^2}$	xy	$0.5(x^2+y^2)$
RS pooling	6.69 ± 6.07	8.43 ± 3.47	4.70 ± 4.53
Graph coarsening	6.13 ± 5.01	7.97 ± 4.25	9.30 ± 4.06
Selection GNN	17.12 ± 12.67	13.07 ± 9.64	13.10 ± 7.40

(b) Source localization with graphs initiated with Regular Sampling (RS) graphon pooling.

Architecture	$W(x,y)$		
	$e^{-\beta(x-y)^2}$	xy	$0.5(x^2+y^2)$
M2 pooling	21.79 ± 11.18	10.75 ± 7.43	23.30 ± 2.50
Graph coarsening	7.81 ± 4.49	4.69 ± 3.69	6.40 ± 5.60
Selection GNN	36.75 ± 11.87	8.62 ± 6.61	26.90 ± 6.51

(c) Source localization with graphs initiated with **M2** graphon pooling.

Architecture	$W(x,y)$		
	$e^{-\beta(x-y)^2})$	xy	$0.5(x^2+y^2)$
M3 pooling	19.81 ± 11.02	11.50 ± 8.69	17.69 ± 7.17
Graph coarsening	9.31 ± 7.87	9.83 ± 9.74	12.31 ± 7.02
Selection GNN	34.50 ± 18.20	11.33 ± 8.73	16.56 ± 7.78

(d) Source localization with graphs initiated with **M3** graphon pooling.

Table I: Source localization test error rates (%) achieved by four graphon pooling methods, graph coarsening and selection GNN on 100-node graphs obtained from exponential, bilinear and polynomial graphons.

We can observe from the results that graphon pooling and graph coarsening outperform selection GNN under an exponential graphon function. With bilinear or polynomial graphon functions, **M2** and **M3**) graphon pooling methods perform worse because the randomness involved ignores the form of the graphon functions. Furthermore, the Regular Sampling and **M1** graphon pooling methods can match or outperform graph coarsening method under all three types of graphon functions. Though the graph coarsening method can achieve a stable performance, it still has limitations in calculation complexity and fixed dimensionality ratio.

We further investigate the influence of the layer-wise dimensionality reduction ratios N_1/N (layer 1) and N_2/N_1 (layer 2) on the performance. We focus on the polynomial graphon $W(x,y) = 0.5(x^2+y^2)$ with different number of nodes N and the numbers of selected nodes N_1 and N_2 as shown in the rows of Table II. Considering the fixed sampling ratios in graph coarsening, we only compare different graphon pooling methods with selection GNN. This also shows that our graphon pooling methods also perform well with reduced complexity, especially when the reduction ratio is large.

B. Point cloud classification problem

We next evaluate the pooling strategies on the Model-Net10 [35] to classify a certain object. The ModelNet10 dataset contains 3,991 meshed CAD models from 10 categories for training and 908 models for testing. We sample 300

$[N, N_1, N_2]$	$\left[\frac{N}{N_1}, \frac{N_1}{N_2}\right]$	**M1** pooling	Selection GNN
$[100, 50, 10]$	$[2, 5]$	5.88 ± 4.71	4.56 ± 3.9
$[200, 100, 10]$	$[2, 10]$	23.40 ± 12.88	25.70 ± 8.33
$[400, 200, 10]$	$[2, 20]$	28.94 ± 9.63	38.81 ± 16.41
$[100, 20, 10]$	$[5, 2]$	10.31 ± 7.36	21.25 ± 10.79
$[200, 20, 10]$	$[10, 2]$	29.56 ± 12.74	44.19 ± 14.74
$[400, 20, 10]$	$[20, 2]$	46.05 ± 14.95	54.70 ± 12.58

(a) Source localization with graphs initiated with **M1** graphon pooling.

$[N, N_1, N_2]$	$\left[\frac{N}{N_1}, \frac{N_1}{N_2}\right]$	RS pooling	Selection GNN
$[100, 50, 10]$	$[2, 5]$	4.75 ± 3.73	3.06 ± 2.54
$[200, 100, 10]$	$[2, 10]$	21.38 ± 8.01	23.25 ± 18.28
$[400, 200, 10]$	$[2, 20]$	24.31 ± 11.81	27.93 ± 12.90
$[100, 20, 10]$	$[5, 2]$	7.06 ± 2.54	9.56 ± 6.40
$[200, 20, 10]$	$[10, 2]$	20.45 ± 6.52	27.85 ± 12.20
$[400, 20, 10]$	$[20, 2]$	42.95 ± 9.23	48.20 ± 11.25

(b) Source localization with graphs initiated with Regular Sampling (RS) graphon pooling.

Table II: Source localization test error rates (%) achieved by graphon pooling, graph coarsening and selection GNN on graphs obtained from $W(x,y) = 0.5(x^2+y^2)$ with different values of N, N_1 and N_2.
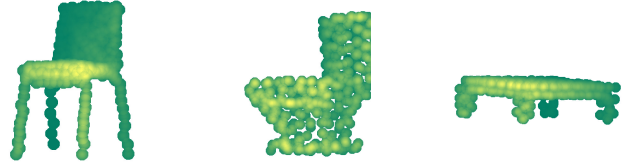

Figure 4. Point cloud models with 300 sampling points in each model. Our goal is to identify chair models from other models such as toilet and table.

points from each model to construct a point cloud. Each point possesses a 3D coordinate as features. We model the graph by seeing the sampling points as nodes and the distance between every pair of nodes as edge weight between two nodes. Based on this graph, we can generate a step graphon function with 300×300 blocks. Our goal is to identify the models for chairs from all the other categories.

We implement graphon pooling, graph coarsening and selection GNN architectures. All these architectures include GNNs with 2 layers with $F_1 = 64$ output features in the first layer and $F_2 = 32$ output features in the second layer. Each layer contains $K = 5$ filter taps. We use ReLU as the activation function. All architectures also include a final readout layer to map the graph output features to a binary classfication scalar. All the architectures are trained by minimizing the cross-entropy loss with the learning rate set as 0.005. We divide the training models into batches of 10 models over 40 epochs. We repeat 5 sampling realizations for all the architectures and evaluate the performance by averaging the classification error rates as well as the standard deviation in Table III.

We can observe that **M1** and Regular Sampling graphon pooling methods outperform graph coarsening and selection

Architecture	$[N_1, N_2]$		
	$[150, 75]$	$[100, 50]$	$[50, 10]$
M1 pooling	2.97 ± 0.23	3.96 ± 1.56	5.14 ± 0.26
RS pooling	2.94 ± 0.34	4.41 ± 1.41	5.80 ± 0.55
M2 pooling	3.23 ± 0.26	4.81 ± 1.54	7.45 ± 1.90
M3 pooling	3.16 ± 0.23	4.44 ± 1.41	7.20 ± 2.71
Selection GNN	2.71 ± 0.27	4.74 ± 1.40	6.64 ± 2.14
Graph Coarsening	2.46 ± 0.46	–	–

Table III: Prediction error rates (%) for model 'chair' in the test dataset. Average over 5 data realizations. The number of nodes is $N = 300$, and $[N_1, N_2]$ stands for the number of selected nodes in the first and second layers of the GNN.

Architecture	$[N_1, N_2] = [150, 75]$	$[N_1, N_2] = [50, 10]$
M1 pooling	$0.405s \pm 0.011s$	$0.067s \pm 0.007s$
RS pooling	$0.083s \pm 0.007s$	$0.028s \pm 0.003s$
M2 pooling	$0.422s \pm 0.018s$	$0.070s \pm 0.009s$
M3 pooling	$0.083s \pm 0.008s$	$0.028s \pm 0.004s$
Selection GNN	$1.936s \pm 0.062s$	$1.101s \pm 0.037s$
Graph coarsening	$5.110s \pm 0.044s$	–

Table IV: Average training time spent per batch in point cloud classification problem. The number of nodes is $N = 300$, and $[N_1, N_2]$ stands for the number of selected nodes in the first and second layers of the GNN.

GNN when the dimensionality ratio is large as the last two columns show in Figure III. We further compare the computation complexity by demonstrating the training time per batch spent in each strategy as Table IV shows. We claim that though graphon pooling may achieve similar results or slightly worse result when compared with graph coarsening, it needs much less time for training. This indicates the graphon pooling method can achieve a better performance with high computation efficiency.

C. Recommendation system problem

We implement the MovieLens 100k dataset [36] to construct a user similarity network. The MovieLens dataset contains 100,000 ratings given from 943 users to 1,682 movies. By calculating Pearson correlations between the ratings given by two users to the same movies [37] while keeping the number of nearest neighbors of each user as 50, we can build a full similarity network. A step graphon function with 943×943 blocks can be generated based on this full network.

The graph signal represents the movies' rating vectors with the u-th element of the rating vector for movie m standing for the rating given by user u to movie m. Given an incomplete rating vector of a specific movie, we can predict a user's rating based on the similarity network.

We train GNNs with graphon pooling and selection GNN architectures by minimizing the mean squared error (MSE) loss between the real and the predicted ratings. All architectures contain 2 layers with $F_1 = 32$ and $F_2 = 8$ features respectively. Each layer consists of $K_1 = K_2 = 5$ filter taps. We focus on user 1 and divide 90% of movie ratings from user 1 for training while the rest for testing. We train all the architectures over 40 epochs with batch size 5.

We set the number of selected nodes as $N_1 = 100$ and $N_2 = 10$ nodes in the first and second layers respectively. We

Architecture	$[N_1, N_2] = [100, 10]$	$[N_1, N_2] = [50, 10]$
M1 pooling	1.1066 ± 0.1497	1.1678 ± 0.1158
RS pooling	1.1121 ± 0.1462	1.1797 ± 0.1358
M2 pooling	1.1558 ± 0.1342	1.1998 ± 0.1380
M3 pooling	1.1132 ± 0.1471	1.2160 ± 0.1301
Selection GNN	1.1339 ± 0.1152	1.2167 ± 0.1006

Table V: Prediction RMSE for user 1's ratings to movies in the test set. Average over 10 train-test splits. The number of nodes is $N = 943$, and $[N_1, N_2]$ stands for the number of selected nodes in the first and second layers of the GNN.

also consider another setting with $N_1 = 50$ and $N_2 = 10$. The average prediction RMSEs and the standard deviations over 10 different data realizations are presented in Table V. For $[N_1, N_2] = [100, 10]$, the GNN with graphon pooling achieves lower test RMSE than the selection GNNs especially the **M1** and Regular Sampling methods, which is accordant with our conclusions in Section V-A and V-B.

VI. DISCUSSION AND CONCLUSIONS

We proposed three pooling strategies for signals and operators on graphs based on graphon representations – Section IV, methods **M1**, **M2**, and **M3**. The underlying idea of **M1/M2/M3** relies on building sequences of graphs and graph signals from graphon representations. We tested our approach on GNNs, making a comparison with other graph pooling approaches such as graph coarsening and zero padding – Section IV.

We proved that methods **M1/M2** based on integration over partitions of $[0, 1]^2$ lead to bounded errors for the filters and mapping operators when compared to a Gphon-NN built only with the original graph/graphon – Theorems 2, 3, and 5. In Theorem 2 we showed that the shift operators on the induced graphons of the graphs obtained after applying **M1/M2** are close to the shift operator on the induced graphon of the original graph. This closeness is inversely proportional to the squared number of elements of the partition considered. In Theorem 3 we showed that the size of the difference between the filters on the resultant induced graphon after applying **M1/M2** and the induced graphon of the original graph, is bounded by the squared root of the cut norm distance between graphon representations. This bound can be calculated using graphs/graphons obtained after applying **M1/M2**. Theorem 4 extends Theorem 3 to the mapping operators of the Gphon-NNs, showing that pointwise nonlinearities do not alter the error associated to the pooling operation. Additionally, we formally showed that the method **M1** – integration over an equipartition of $[0, 1]$ – is stable to arbitrary approximations of the graph and its induced graphon representation. This is, given two graph approximations of a graphon that are ϵ-close, we have that **M1** preserves the closeness between such approximations up to a scalar factor – Theorem 5.

The numerical experiments in Section V corroborate the results derived from our analysis in Section IV, showing that among the graphon pooling methods proposed **M1** provides the best results. This points to a unique attribute of equipartitions in the graphon space, which is partly explained by the stability of **M1** to arbitrary approximations of the graph to be reduced – Theorem 5.

When applied to reduce dimensionality between layers in GNNs, graphon pooling methods adapt particularly well to those scenarios where the graph in the first layer is large. This is a consequence of two fundamental facts. First, graphons naturally model graphs of large size. Second, the graphon pooling methods are more effective than other pooling approaches when there is a large dimensionality reduction ratio between the original signal/operator and the reduced one.

Although valuable for medium size graphs, pooling methods such as graph coarsening and zero padding are not applicable for large graphs because of their computational cost. In contrast with this, graphon pooling adapts naturally to large graphs since graphons are by themselves limit objects of sequences of graphs whose number of nodes and edges grows up to infinity.

APPENDIX A
BACKGROUND MATERIAL

A. Basic preliminaries

1) Spectral representation of convolutional filtering on graphs: If $\mathbf{S}$ is diagonalizable, we can leverage the spectral representation of $\mathbf{S}$ to define a Fourier transform and a spectral representation of the filters. Let $\mathbf{S} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{\mathsf{T}}$, where $\mathbf{\Lambda}$ is diagonal and $\mathbf{U}$ is orthogonal. Then, we say that $\hat{\mathbf{x}} = \mathbf{U}^{\mathsf{T}}\mathbf{x}$ is the graph Fourier transform (GFT) of $\mathbf{x}$. The filtering of $\mathbf{x}$ can be expressed in the Fourier domain according to $\mathbf{H}(\mathbf{S})\mathbf{x} = \mathbf{U}\sum_{k=0}^{K-1} h_k \mathbf{\Lambda}^k \hat{\mathbf{x}}$. This is, to do the filtering of $\mathbf{x}$ by $\mathbf{H}(\mathbf{S})$ we can do the spectral filtering of $\hat{\mathbf{x}}$ as $\sum_{k=0}^{K-1} h_k \mathbf{\Lambda}^k \hat{\mathbf{x}}$, and then take the inverse Fourier transform given by the right action of $\mathbf{U}$.

2) Spectral representation of convolutional filtering on graphons: Using spectral decompositions of $\boldsymbol{T}_W$, it is possible to define a Fourier transform on graphons. Suppose $\lambda_i(\boldsymbol{T}_W)$ and $\boldsymbol{\varphi}_{W,i}$ are the i-th eigenvalue and i-th eigenvector of $\boldsymbol{T}_W$, respectively. We denote the graphon Fourier transform (Gphon-FT) of $(W, \boldsymbol{x})$ by $(\widehat{W}, \hat{\boldsymbol{x}})$, where $\hat{\boldsymbol{x}} \in \ell_2(\mathbb{Z})$ with

$$\hat{\boldsymbol{x}}(j) = \int_0^1 \boldsymbol{x}(u)\boldsymbol{\varphi}_{W,j}(u)du, \qquad (19)$$

and where the symbol $\widehat{W}$ emphasizes that $\hat{\boldsymbol{x}}$ is defined on the spectrum – eigenvalues – of $\boldsymbol{T}_W$ and not in $[0, 1]$.

There is a natural connection between the GFT and the Gphon-FT when a graphon signal is induced by a graph signal. If W_G is induced by G, we have $\lambda_i(\boldsymbol{T}_{W_G}) = \lambda_i(G)/|V(G)|$ for $i = 1, \ldots, |V(G)|$ and $\lambda_i(\boldsymbol{T}_{W_G}) = 0$ for all $i > |V(G)|$ [26], [30], considering an ordering of the eigenvalues in decreasing absolute modulus. Furthermore, if $\mathbf{u}_i$ is the i-th eigenvector of $\mathbf{S}$ – the shift operator of the graph G –, the i-th eigenvector of $\boldsymbol{T}_{W_G}$ is given by $\boldsymbol{\varphi}_{W_G,i}(t) = \sqrt{|V(G)|}\mathbf{u}_i(\lceil t|V(G)|\rceil)$ with $N = |V(G)|$ [38].

The convergence of a sequence of graphs to a graphon has implications on the convergence of spectral representations on the graphs and the graphons [26]. This is formally stated in the following theorem from [26].

Theorem 6 (5.6 [26]). *Let $\{W_i\}_i$ be a sequence of graphons uniformly bounded in L^∞. Suppose $\|W_i - W\|_\square \to 0$ as $i \to \infty$. For $k \geq 1$, let $P_k(W_i): L^2([0, 1]) \to L^2([0, 1])$ be*

the projection operator on the eigenspace of W_i associated to the eigenvalues $\{\lambda_k(W_i), -\lambda_k(W_i)\}$. For all $k \geq 1$ we have

1) $\lambda_k(W_i) \to \lambda_k(W)$ *as $i \to \infty$,*
2) $\|P_k(W_i) - P_k(W)\|_2 \to 0$ *as $i \to \infty$.*

3) Convergence of homomorphism densities: In this subsection we discuss the notion of convergence of sequences of graphs to graphs, relying on the notion of homomorphisms density. A homomorphism from the graph H to the graph G is an edge preserving map from $V(H)$ to $V(G)$. If we denote by $\mathrm{Hom}(H, G)$ the number of homomorphisms between H and G, it is possible to define a homomorphism density $t(H, G)$ by $t(H, G) = |\mathrm{Hom}(H, G)|/|V(G)|^{|V(H)|}$. This notion of density can be extended to compare graphs and graphons. If H is a graph and $W(x, y)$ is a graphon, the density of homomorphisms between H and $W(x, y)$ can be calculated as

$$t(H, W) = \int_{[0,1]^{|V(H)|}} \prod_{\{i,j\}\in E(H)} W(x_i, x_j)dx_1 \ldots dx_{|V(H)|}. \qquad (20)$$

If W_G is a graphon induced by a graph G, then it is possible to show that $t(H, W_G) = t(H, G)$ [26], [30], [31].

We say that a sequence of graphs $\{G_i\}_{i=0}^\infty$ converge to the graphon $W(x, y)$ – we denote this by $\{G_i\} \to W(x, y)$ – if $\lim_{i\to\infty} t(H, G_i) = t(H, W)$ for all simple finite graphs H.

APPENDIX B
PROOFS

We start introducing some notation that will facilitate the presentation of the proofs. The symbol $\mathcal{B}(\mathcal{V})$ indicates the set of bounded operators acting on $\mathcal{V}$. The norm indicated as $\|\cdot\|_2$ represents the norm $\|\cdot\|_{2\to 2}$ on operators in $\mathcal{B}(L^2([0,1]))$ – the set of bounded operators acting on $L^2([0, 1])$. The norm indicated as $\|\cdot\|_{op2}$ represents the norm $\|\cdot\|_{\mathcal{B}(L^2([0,1]))\to\mathcal{B}(L^2([0,1]))}$ acting on elements of $\mathcal{B}(\mathcal{B}(L^2([0,1])))$. The symbol $\|\cdot\|_{op(\mathcal{B}\otimes\mathcal{B})}$ represents the operator norm of an operator living in $\mathcal{B}(L^2([0,1]))^* \otimes \mathcal{B}(L^2([0,1]))$. The symbol $\|\cdot\|_{L^2\otimes L^2}$ is the norm acting on $L^2([0,1]) \otimes L^2([0,1])$. Given the graphon $W(x, y)$ we will denote by $\|W\|_1$ and $\|W\|_2$ the L^1 and L^2 norms of $W(x, y)$ as elements of $L^1([0,1]^2)$ and $L^2([0,1]^2)$.

A. Proof of Theorem 1

First, we start taking into account that by means of the spectral theorem we have

$$\boldsymbol{y} = h(\boldsymbol{T}_{W_G})\boldsymbol{x} = \sum_{i=1}^\infty h(\lambda_i(\boldsymbol{T}_{W_G}))\boldsymbol{\varphi}_{W_G,i}\langle\boldsymbol{\varphi}_{W_G,i}, \boldsymbol{x}\rangle. \qquad (21)$$

Taking into account that

$$\boldsymbol{\varphi}_{W_G,i} = \sqrt{|V(G)|}\sum_{r=1}^{|V(G)|}\mathbf{u}_i(r)\chi_r(u), \quad \boldsymbol{x} = \sum_{r=1}^{|V(G)|}\mathbf{x}(r)\chi_r(u), \qquad (22)$$

where $\mathbf{u}_i$ is the i-th eigenvector of $\mathbf{S}_G$ and $\chi_r(u)$ is the characteristic function of the interval $[(r-1)/|V(G)|, r/|V(G)|]$, it follows that

$$
\begin{aligned}
\langle \varphi_{W_G,i}, \mathbf{x} \rangle &= \\
\int_0^1 \sqrt{|V(G)|} &\sum_{r=1}^{|V(G)|} \mathbf{u}_i(r)\chi_r(u) \sum_{r=1}^N \mathbf{x}(r)\chi_r(u)du \\
&= \frac{\langle \mathbf{u}_i, \mathbf{x} \rangle}{\sqrt{|V(G)|}}. \quad (23)
\end{aligned}
$$

Then, taking into account that $\lambda_i(\boldsymbol{T}_{W_G}) = \lambda_i(G)/|V(G)|$ and $\varphi_{W_G,i} = \sqrt{|V(G)|}\text{step}(\mathbf{u}_i)$ we reach

$$
\boldsymbol{y} = \sum_{i=1}^\infty h\left(\frac{\lambda_i(G)}{|V(G)|}\right) \sqrt{|V(G)|}\text{step}(\mathbf{u}_i) \frac{\langle \mathbf{u}_i, \mathbf{x} \rangle}{\sqrt{|V(G)|}}. \quad (24)
$$

Using the properties of $\text{step}(\cdot)$ we have

$$
\boldsymbol{y} = \text{step}\left(\sum_{i=1}^\infty h\left(\frac{\lambda_i(G)}{|V(G)|}\right) \mathbf{u}_i \langle \mathbf{u}_i, \mathbf{x} \rangle\right), \quad (25)
$$

and by means of the spectral theorem it follows that

$$
\boldsymbol{y} = \text{step}\left(h\left(\frac{\mathbf{S}_G}{|V(G)|}\right) \mathbf{x}\right). \quad (26)
$$

B. Proof of Theorem 2

Proof. This proof follows directly from the application of Proposition 9.8 in [30]. To see this, we start recalling the notation in [30] where for a given graphon $W(x,y)$ its *stepping* is given by

$$
W_{\mathcal{P}}(x,y) = \frac{1}{\mu(\mathcal{S}_i)\mu(\mathcal{S}_j)} \int_{\mathcal{S}_i \times \mathcal{S}_j} W(x,y)dxdy, \quad (27)
$$

where μ is a measure function and $\mathcal{P} = \{\mathcal{S}_i\}_i$ is a partition of $[0,1]$. Now, we recall Proposition 9.8 in [30].

Theorem 7 (Proposition 9.8 [30]). *Let $\{\mathcal{P}_i\}_i$ be a sequence of measureable partitions of $[0,1]$ such that every pair of points is separated by all but a finite number of partitions $\mathcal{P}_i$. Then, $W_{\mathcal{P}_i} \to W(x,y)$ almost everywhere for every $W \in \mathcal{W}$, where $\mathcal{W}$ is the space of kernels on $[0,1]^2$.*

The graphs built from methods **M1** and **M2** are steppings of the graphon W. Additionally, we remark that every graph obtained using **M1** and **M2** is associated to nested partitions of $[0,1]$. This is, if $I_\ell^{(i)}$ are the intervals in a partition of $[0,1]$ associated to $\mathcal{P}_i$ and $I_\ell^{(j)}$ are the intervals in a partition associated to $\mathcal{P}_j$, we have that $I_\ell^{(i)} = \bigcup_{\ell \in \mathcal{J}} I_\ell^{(j)}$, where $\mathcal{J} \subset \mathbb{N}$. Then, the condition stipulated in Theorem 7 is satisfied and therefore, the sequence of graphons $\{W_{G_i}\}_i$ converges almost everywhere to W.

Now, from Lemma 9.11 in [30] we have

$$
\|W - W_{G_\ell}\|_{\square} \leq \frac{2}{\sqrt{\log |V(G_\ell)|}}. \quad (28)
$$

Taking into account that $\|\boldsymbol{T}_W\|_{\infty \to 1} \leq 4\|W\|_{\square}$ – see [30] page 134 – it follows that

$$
\left\|\boldsymbol{T}_W - \boldsymbol{T}_{W_{G_\ell}}\right\|_{\infty \to 1} \leq \frac{8}{\sqrt{\log |V(G_\ell)|}}. \quad (29)
$$

$\square$

C. Fréchet derivative of a Graph/Graphon Filter

In this section we show the details for the calculation of the Fréchet derivative of a graph filter instantiated in the induced graphon representation.

Theorem 8. *Let $\boldsymbol{T}_W$ be the graphon shift operator of the graphon W, and let $h(\boldsymbol{T}_W) = \sum_{k=0}^\infty h_k \boldsymbol{T}_W^k$. If $D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\}$ is the Fréchet derivative of $h(\boldsymbol{T}_W)$, acting on $\boldsymbol{\xi}$, it follows that*

$$
D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\} = \sum_{i=1}^\infty \ell_i(\boldsymbol{T}_W)\boldsymbol{\xi} r_i(\boldsymbol{T}_W), \quad (30)
$$

where $\ell_i(\boldsymbol{T}_W)$ and $r_i(\boldsymbol{T}_W)$ are monomial functions.

Proof. Let us start taking into account that

$$
h(\boldsymbol{T}_W + \boldsymbol{\xi}) - h(\boldsymbol{T}_W) = \sum_{k=0}^\infty h_k (\boldsymbol{T}_W + \boldsymbol{\xi})^k - \sum_{k=0}^\infty h_k \boldsymbol{T}_W^k. \quad (31)
$$

We expand the polynomial expression and re-group the monomials in two terms. One term where $\boldsymbol{\xi}$ pears once and another term where $\boldsymbol{\xi}$ appears more than once. Then, we have

$$
h(\boldsymbol{T}_W + \boldsymbol{\xi}) - h(\boldsymbol{T}_W) = \sum_{k=0}^\infty h_k \sum_{\ell=1}^k \boldsymbol{T}_W^{\ell-1}\boldsymbol{\xi}\boldsymbol{T}_W^{k-\ell} + o\left(\|\boldsymbol{\xi}\|_2\right), \quad (32)
$$

with $o\left(\|\boldsymbol{\xi}\|_2\right) \to 0$ as $\|\boldsymbol{\xi}\|_2 \to 0$. Notice that the term $o\left(\|\boldsymbol{\xi}\|_2\right)$ is a polynomial function where $\boldsymbol{\xi}$ appears more than once in each monomial. Since $\sum_{k=0}^\infty h_k \sum_{\ell=1}^k \boldsymbol{T}_W^{\ell-1}\boldsymbol{\xi}\boldsymbol{T}_W^{k-\ell}$ is linear and bounded as an operator on $\boldsymbol{\xi}$ we have from the definition of the Fréchet derivative [39], [40] that

$$
D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\} = \sum_{k=0}^\infty h_k \sum_{\ell=1}^k \boldsymbol{T}_W^{\ell-1}\boldsymbol{\xi}\boldsymbol{T}_W^{k-\ell}, \quad (33)
$$

which is indeed a sum of monomials in $\boldsymbol{T}_W$ acting on the left and right of $\boldsymbol{\xi}$. $\square$

D. Proof of Theorem 3

Let $\boldsymbol{\xi} = \boldsymbol{T}_W - \boldsymbol{T}_{W_G}$ and $D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\}$ be the Fréchet derivative of $h : \mathcal{B}(L^2[0,1]) \to \mathcal{B}(L^2[0,1])$. Then, from the definition of Fréchet derivative – see proof of Theorem 8 in Section B-C – we have

$$
h(\boldsymbol{T}_W) - h(\boldsymbol{T}_{W_G}) = D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\} + o\left(\|\boldsymbol{\xi}\|_2\right), \quad (34)
$$

where $D_{h|\boldsymbol{T}_W}\{\cdot\}$ is linear and bounded. Taking the norm on both sides of the equation above and using the triangle inequality we have

$$
\|h(\boldsymbol{T}_W) - h(\boldsymbol{T}_{W_G})\|_2 \leq \left\|D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\}\right\|_2 + \mathcal{O}\left(\|\boldsymbol{\xi}\|_2^2\right). \quad (35)
$$

Now, we focus our attention on the term $D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\}$. Since $h(\cdot)$ is a polynomial function, from Theorem 8 we can express $D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\}$ as $D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\} = \sum_{i=1}^{\infty} \ell_i(\boldsymbol{T}_W)\boldsymbol{\xi} r_i(\boldsymbol{T}_W)$, where ℓ_i and r_i are polynomial functions. Since $\boldsymbol{T}_W$ is Hilbert-Schmidt, $D_{h|\boldsymbol{T}_W}\{\cdot\}$ is also Hilbert-Schmidt. Then, as stated in [41] (page 268) there is an isomorphic-isometric image of $D_{h|\boldsymbol{T}_W}\{\cdot\} \in \mathcal{B}\left(\mathcal{B}\left(L^2[0,1]\right)\right)$ in $\mathcal{B}(L^2[0,1])^* \otimes \mathcal{B}(L^2[0,1])$ that we denote by $\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}$. From [42] (page 61), the eigenvalues of $\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}$ can be written as

$$\lambda_{i,j}\left(\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}\right) =$$

$$\begin{cases} \frac{h(\lambda_i(\boldsymbol{T}_W)) - h(\lambda_j(\boldsymbol{T}_W))}{\lambda_i(\boldsymbol{T}_W) - \lambda_j(\boldsymbol{T}_W)}, & \lambda_i(\boldsymbol{T}_W) \neq \lambda_j(\boldsymbol{T}_W) \\ h'\left(\lambda_i(\boldsymbol{T}_W)\right), & \lambda_i(\boldsymbol{T}_W) = \lambda_j(\boldsymbol{T}_W), \end{cases} \quad (36)$$

while the eigenvectors of $\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}$ are given by $\boldsymbol{\varphi}_{i,j} = \boldsymbol{\varphi}_i \otimes \boldsymbol{\varphi}_j$. If $\boldsymbol{\xi} \in \mathcal{B}(L^2[0,1])$ is Hilbert-Schmidt, there exists $\overline{\boldsymbol{\xi}} \in L^2[0,1] \otimes L^2[0,1]$ defined by the isomorphic-isometric map between $\mathcal{B}\left(\mathcal{B}\left(L^2[0,1]\right)\right)$ and $\mathcal{B}(L^2[0,1])^* \otimes \mathcal{B}(L^2[0,1])$.

With these facts at hand we start taking into account that from the operator norm it follows

$$\left\|D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\}\right\|_2 \leq \left\|D_{h|\boldsymbol{T}_W}\{\cdot\}\right\|_{op2} \|\boldsymbol{\xi}\|_2. \quad (37)$$

Then, since $\left\|D_{h|\boldsymbol{T}_W}\{\cdot\}\right\|_{op2} = \left\|\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}\right\|_{op(\mathcal{B}\otimes\mathcal{B})}$, and $\|\boldsymbol{\xi}\|_2 \leq \|\boldsymbol{\xi}\|_{HS} = \|\overline{\boldsymbol{\xi}}\|_{L^2 \otimes L^2}$, we have

$$\left\|D_{h|\boldsymbol{T}_W}\{\boldsymbol{\xi}\}\right\|_2 \leq \left\|\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}\right\|_{op(\mathcal{B}\otimes\mathcal{B})} \|\boldsymbol{\xi}\|_{HS}. \quad (38)$$

Replacing (38) in (35) we have

$$\|h(\boldsymbol{T}_W) - h(\boldsymbol{T}_{W_G})\|_2 \leq \\ \left\|\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}\right\|_{op(\mathcal{B}\otimes\mathcal{B})} \|\boldsymbol{\xi}\|_{HS} + \mathcal{O}\left(\|\boldsymbol{\xi}\|_2^2\right). \quad (39)$$

Since $h(t)$ is L-Lipschitz, from (36) we have that $\left\|\overline{D}_{h|\boldsymbol{T}_W}\{\cdot\}\right\|_{op(\mathcal{B}\otimes\mathcal{B})} \leq L$, and therefore (39) turns into

$$\|h(\boldsymbol{T}_W) - h(\boldsymbol{T}_{W_G})\|_2 \leq L \|\boldsymbol{\xi}\|_{HS} + \mathcal{O}\left(\|\boldsymbol{\xi}\|_2^2\right). \quad (40)$$

Since $\|\boldsymbol{\xi}\|_{HS} \leq \gamma \|\boldsymbol{\xi}\|_2$ and $\|\boldsymbol{T}_W\|_2 \leq \sqrt{8\|W\|_\square}$ (see Proposition 4, page 15 in [28]) we have

$$\|h(\boldsymbol{T}_W) - h(\boldsymbol{T}_{W_G})\|_2 \leq L\gamma\sqrt{8\|W - W_G\|_\square} + \mathcal{O}\left(\|\boldsymbol{\xi}\|_2^2\right). \quad (41)$$

Taking into account Lemma 1, there exists a graph H with $V(H) \leq 4V(G)$ such that

$$\|h(\boldsymbol{T}_W) - h(\boldsymbol{T}_{W_G})\|_2 \leq L\gamma\sqrt{8\|W_H - W_G\|_\square} + \mathcal{O}\left(\|\boldsymbol{\xi}\|_2^2\right). \quad (42)$$

E. Proof of Theorem 4

Proof. First, we are going to estimate the upper bound of the difference between the perceptron operators in each layer of the graphon neural networks. Since η_ℓ is 1-Lipschitz we have

$$\left\|\eta_\ell\left(h_\ell(\boldsymbol{T}_W)\boldsymbol{x}_{\ell-1}\right) - \eta_\ell\left(h_\ell(\boldsymbol{T}_{W_{G_\ell}})\boldsymbol{x}_{\ell-1}\right)\right\|_2 \\ \leq \left\|h_\ell(\boldsymbol{T}_W)\boldsymbol{x}_{\ell-1} - h_\ell(\boldsymbol{T}_{W_{G_\ell}})\boldsymbol{x}_{\ell-1}\right\|_2. \quad (43)$$

Now, by means of the operator norm property we have

$$\left\|\eta_\ell\left(h_\ell(\boldsymbol{T}_W)\boldsymbol{x}_{\ell-1}\right) - \eta_\ell\left(h_\ell(\boldsymbol{T}_{W_{G_\ell}})\boldsymbol{x}_{\ell-1}\right)\right\|_2 \\ \leq \left\|h_\ell(\boldsymbol{T}_W) - h_\ell(\boldsymbol{T}_{W_{G_\ell}})\right\|_2 \|\boldsymbol{x}_{\ell-1}\|_2. \quad (44)$$

In what follows we will use the notation $\boldsymbol{E}_\ell = \left\|h_\ell(\boldsymbol{T}_W) - h_\ell(\boldsymbol{T}_{W_{G_\ell}})\right\|_2$.

Now, we analyze the difference between the output signals in the ℓ-th layer, which we can write as follows

$$\|\boldsymbol{x}_\ell - \tilde{\boldsymbol{x}}_\ell\|_2 \leq \|\eta_{\ell-1}h_{\ell-1}(\boldsymbol{T}_W)\eta_{\ell-2}h_{\ell-2}(\boldsymbol{T}_W)\cdots \\ \eta_1 h_1(\boldsymbol{T}_W)\boldsymbol{x} - \\ \eta_{\ell-1}h_{\ell-1}(\boldsymbol{T}_{W_{G_{\ell-1}}})\eta_{\ell-2}h_{\ell-2}(\boldsymbol{T}_{W_{G_{\ell-2}}})\cdots \\ \eta_1 h_1(\boldsymbol{T}_{W_{G_1}})\boldsymbol{x}\|_2. \quad (45)$$

Then, we take into account that

$$h_{\ell+1}(\boldsymbol{T}_W)\eta_\ell(a) - h_{\ell+1}(\boldsymbol{T}_{W_{G_{\ell+1}}})\eta_\ell(\tilde{a}) = \\ (h_{\ell+1}(\boldsymbol{T}_W) - h_{\ell+1}(\boldsymbol{T}_{W_{G_{\ell+1}}}))\eta_\ell(a) + \\ h_{\ell+1}(\boldsymbol{T}_{W_{G_\ell}})(\eta_\ell(a) - \eta_\ell(\tilde{a})), \quad (46)$$

where a and $\tilde{a}$ indicate the terms that are on the right side of η_ℓ in (46). Since $\|\eta_\ell(a) - \eta_\ell(b)\|_2 \leq \|a - b\|_2$, $\|h_{\ell+1}(\boldsymbol{T}_W) - h_{\ell+1}(\boldsymbol{T}_{W_{G_{\ell+1}}})\|_2 \leq \boldsymbol{E}_{\ell+1}$, and $\|h_{\ell+1}\|_2 \leq 1$, we have that

$$\left\|h_{\ell+1}(\boldsymbol{T}_W)\eta_\ell(a) - h_{\ell+1}(\boldsymbol{T}_{W_{G_{\ell+1}}})\eta_\ell(\tilde{a})\right\|_2 \leq \\ \boldsymbol{E}_{\ell+1}\|a\|_2 + \|a - \tilde{a}\|_2. \quad (47)$$

Taking into account these results recursively on the index ℓ we have

$$\left\|\Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\mathcal{S}_\ell\}_{\ell=1}^L\right) - \Phi\left(\boldsymbol{x}, \{\mathcal{P}_\ell\}_{\ell=1}^L, \{\tilde{\mathcal{S}}_\ell\}_{\ell=1}^L\right)\right\|_2 \\ \leq \sum_{\ell=1}^L \boldsymbol{E}_\ell \|\boldsymbol{x}\|_2. \quad (48)$$

Finally, we take into account that $\boldsymbol{E}_\ell = \left\|h_\ell(\boldsymbol{T}_W) - h_\ell(\boldsymbol{T}_{W_{G_\ell}})\right\|_2$ is bounded according to Theorem 3. This completes the proof. $\square$

F. Proof of Theorem 5

Proof. First we take into account that by means of Lemma 5 in [28] (page 16) we have $\|W - W_{G_i}\|_2 = \|\boldsymbol{T}_W - \boldsymbol{T}_{W_{G_i}}\|_{HS}$. Then, taking into account that $\|W\|_1 \leq \|W\|_2$ for any graphon W – see [30], page 131 – it follows that $\|W - W_{G_i}\|_1 \leq \epsilon$.

Now, since $\|W\|_\square \leq \|\boldsymbol{T}_W\|_{\infty \to 1}$ – see [30], page 134 – we have

$$\|W_{G_i} - W_{H_i}\|_\square \leq \frac{\epsilon}{V(H_i)^4}. \quad (49)$$

Then, taking into account Theorem 9.32 in [30] we have

$$\left\| W_{H_1} - W_{H_2}^{\theta} \right\|_1 \leq 8\epsilon. \tag{50}$$

Finally, since $\|\boldsymbol{T}_W\|_{\infty \to 1} \leq 4\|W\|_{\square}$ and $\|W\|_{\square} \leq \|W\|_1$ we have

$$\left\| \boldsymbol{T}_{W_{H_1}} - \boldsymbol{T}_{W_{H_2}^{\theta}} \right\|_{\infty \to 1} \leq 32\epsilon. \tag{51}$$

$\square$

REFERENCES

[1] Alejandro Parada-Mayorga, Luana Ruiz, and Alejandro Ribeiro, "Graphon pooling in graph neural networks," in *2020 28th European Signal Processing Conference (EUSIPCO)*, 2021, pp. 860–864.

[2] Y-Lan Boureau, Jean Ponce, and Yann LeCun, "A theoretical analysis of feature pooling in visual recognition," in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, Madison, WI, USA, 2010, ICML'10, p. 111–118, Omnipress.

[3] Y-Lan Boureau, Nicolas Le Roux, Francis Bach, Jean Ponce, and Yann LeCun, "Ask the locals: Multi-way local pooling for image recognition," in *2011 International Conference on Computer Vision*, 2011, pp. 2651–2658.

[4] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, Adaptive Computation and Machine Learning series. MIT Press, 2016.

[5] Thomas Wiatowski and Helmut Bölcskei, "A mathematical theory of deep convolutional neural networks for feature extraction," *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1845–1866, 2018.

[6] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Velickovi'c, "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges," *ArXiv*, vol. abs/2104.13478, 2021.

[7] A.V. Oppenheim and R.W. Schafer, *Digital Signal Processing*, MIT video course. Prentice-Hall, 1975.

[8] 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. 64, no. 14, pp. 3775–3789, July 2016.

[9] M. Tsitsvero, S. Barbarossa, and P. Di Lorenzo, "Signals on graphs: Uncertainty principle and sampling," *IEEE Transactions on Signal Processing*, vol. 64, no. 18, pp. 4845–4860, Sep. 2016.

[10] P. Di Lorenzo, S. Barbarossa, , and P. Banelli, "Sampling and recovery of graph signals," in *Cooperative and Graph Signal Processing*, P. Djuric and C.Richard, Eds. Elsevier, 2018.

[11] X. Wang, P. Liu, and Y. Gu, "Local-set-based graph signal reconstruction," *IEEE Transactions on Signal Processing*, vol. 63, no. 9, pp. 2432–2444, May 2015.

[12] S. Chen, R. Varma, A. Singh, and J. Kovačević, "Signal recovery on graphs: Fundamental limits of sampling strategies," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 4, pp. 539–554, Dec 2016.

[13] Alejandro Parada-Mayorga, Daniel L. Lau, Jhony H. Giraldo, and Gonzalo R. Arce, "Blue-noise sampling on graphs," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 5, no. 3, pp. 554–569, 2019.

[14] Alejandro Parada-Mayorga, Daniel L. Lau, Jhony H. Giraldo, and Gonzalo R. Arce, "Sampling of graph signals with blue noise dithering," in *2019 IEEE Data Science Workshop (DSW)*, 2019, pp. 150–154.

[15] Alejandro Parada-Mayorga, Daniel L. Lau, Jhony H. Giraldo, and Gonzalo R. Arce, "Blue-noise sampling of signals on graphs," in *2019 13th International conference on Sampling Theory and Applications (SampTA)*, 2019, pp. 1–5.

[16] Dominique Guillot, Alejandro Parada-Mayorga, Sebastian Cioaba, and Gonzalo R. Arce, "Optimal sampling sets in cographs," in *2019 IEEE Data Science Workshop (DSW)*, 2019, pp. 165–169.

[17] Alejandro Parada-Mayorga, *Blue Noise and Optimal Sampling on Graphs*, Ph.D. thesis, 2019.

[18] A. Ortega, *Introduction to Graph Signal Processing*, Cambridge University Press, 2022.

[19] Aliaksei Sandryhaila and José M. F. Moura, "Discrete signal processing on graphs," *IEEE Transactions on Signal Processing*, vol. 61, no. 7, pp. 1644–1656, 2013.

[20] Markus Püschel and José M. F. Moura, "Algebraic signal processing theory," *ArXiv*, vol. abs/cs/0612077, 2006.

[21] Aliaksei Sandryhaila and José M. F. Moura, "Discrete signal processing on graphs: Graph filters," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 6163–6166.

[22] A. Ortega, P. Frossard, J. Kovačević, J. M. F. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 808–828, May 2018.

[23] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, "Convolutional neural network architectures for signals supported on graphs," *IEEE Transactions on Signal Processing*, vol. 67, no. 4, pp. 1034–1049, 2019.

[24] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *NIPS*, 2016.

[25] I. S. Dhillon, Y. Guan, and B. Kulis, "Weighted graph cuts without eigenvectors a multilevel approach," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 11, pp. 1944–1957, 2007.

[26] Peter Diao, Dominique Guillot, Apoorva Khare, and Bala Rajaratnam, "Model-free consistency of graph partitioning," *arXiv: Combinatorics*, 2016.

[27] Matthew W. Morency and Geert Leus, "Graphon filters: Graph signal processing in the limit," *IEEE Transactions on Signal Processing*, vol. 69, pp. 1740–1754, 2020.

[28] Luana Ruiz, Luiz F. O. Chamon, and Alejandro Ribeiro, "Graphon signal processing," 2020.

[29] Fernando Gama, Joan Bruna, and Alejandro Ribeiro, "Stability properties of graph neural networks," 2019.

[30] L. Lovász, *Large Networks and Graph Limits*, American Mathematical Society colloquium publications. American Mathematical Society, 2012.

[31] Daniel Glasscock, "What is a graphon," *arXiv: Combinatorics*, 2015.

[32] Alejandro Parada-Mayorga and Alejandro Ribeiro, "Algebraic neural networks: Stability to deformations," *IEEE Transactions on Signal Processing*, vol. 69, pp. 3351–3366, 2021.

[33] Alejandro Parada-Mayorga and Alejandro Ribeiro, "Stability of algebraic neural networks to small perturbations," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5205–5209.

[34] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[35] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao, "3d shapenets: A deep representation for volumetric shapes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1912–1920.

[36] F Maxwell Harper and Joseph A Konstan, "The movielens datasets: History and context," *Acm transactions on interactive intelligent systems (tiis)*, vol. 5, no. 4, pp. 1–19, 2015.

[37] Weiyu Huang, Antonio G. Marques, and Alejandro R. Ribeiro, "Rating prediction via graph signal processing," *IEEE Transactions on Signal Processing*, vol. 66, no. 19, pp. 5066–5081, 2018.

[38] S. Gao, *Graphon Control Theory for Linear Systems on Complex Networks and Related Topics*, McGill theses. McGill University Libraries, 2019.

[39] Y. Benyamini and J. Lindenstrauss, *Geometric Nonlinear Functional Analysis*, Number v. 48, no. 1 in American Mathematical Society colloquium publications. American Mathematical Society, 2000.

[40] J. Lindenstrauss, D. Preiss, and J. Tišer, *Frechet Differentiability of Lipschitz Functions and Porous Sets in Banach Spaces*, Annals of Mathematics Studies. Princeton University Press, 2012.

[41] J.B. Conway, *A Course in Functional Analysis*, Graduate Texts in Mathematics. Springer New York, 2019.

[42] N.J. Higham, *Functions of Matrices: Theory and Computation*, Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104), 2008.