
On the Universality of Graph Neural Networks on Large Random Graphs

Nicolas Keriven

CNRS, GIPSA-lab, Grenoble, France
nicolas.keriven@cnrs.fr

Alberto Bietti

NYU Center for Data Science, New York, USA
alberto.bietti@nyu.edu

Samuel Vaiter

CNRS, LJAD, Nice, France
samuel.vaiter@cnrs.fr

Abstract

We study the approximation power of Graph Neural Networks (GNNs) on latent position random graphs. In the large graph limit, GNNs are known to converge to certain “continuous” models known as c-GNNs, which directly enables a study of their approximation power on random graph models. In the absence of input node features however, just as GNNs are limited by the Weisfeiler-Lehman isomorphism test, c-GNNs will be severely limited on simple random graph models. For instance, they will fail to distinguish the communities of a well-separated Stochastic Block Model (SBM) with constant degree function. Thus, we consider recently proposed architectures that augment GNNs with unique node identifiers, referred to as Structural GNNs here (SGNNs). We study the convergence of SGNNs to their continuous counterpart (c-SGNNs) in the large random graph limit, under new conditions on the node identifiers. We then show that c-SGNNs are strictly more powerful than c-GNNs in the continuous limit, and prove their universality on several random graph models of interest, including most SBMs and a large class of random geometric graphs. Our results cover both permutation-invariant and permutation-equivariant architectures.

1 Introduction

Graph Neural Networks (GNNs) are deep architectures defined over graph data that have garnered a lot of attention in recent years. They represent the state-of-the-art in many graph Machine Learning (graph ML) problems, and have been successfully applied to e.g. node clustering [7], semi-supervised learning [24], quantum chemistry [17], and so on. See [5, 44, 18, 21] for reviews.

As the universality of Multi-Layers Perceptrons (MLP) is one of the foundational theorems in deep learning – that is, any continuous function can be approximated arbitrarily well by an MLP – in the last few years the approximation power of GNNs has been a topic of great interest. In the absence of special node features, i.e. when one has only access to the graph structure, the crux of the problem has been proven to be the capacity of GNNs to solve the *graph isomorphism problem*, that is, deciding when two graphs are permutations of each other or not (a difficult problem for which no polynomial algorithm is known [4]). Indeed, this property is directly linked to the approximation power of GNNs [8, 3]. In this light, the landmark paper [45] proves that classical GNNs are at best as powerful as the famous Weisfeiler-Lehman (WL) test [43] for graph isomorphism. Since then, several works [45, 31, 8] have derived new architectures, for instance involving high-order tensors [31], with improved discriminative power equivalent to “higher-order” variants of the WL test. In another line of works, several recent papers have advocated the use of *unique node identifiers* [28, 27], with

strategies to preserve the permutation-equivariance/invariance of GNNs coined Structural Message Passing (SMP) in [41] (see Sec. 2.3). We call these models Structural GNN (SGNN) here. SGNNs have been proved to be strictly more powerful than the WL test in [41], and even *universal* on graphs with bounded degrees, however for powerful layers that cannot be implemented in practice.

When the size of the graphs grows, the notion of graph isomorphism becomes somewhat moot: large graphs might share properties (number of well-connected communities, and so on) but are never isomorphic to each other. GNNs have nevertheless proven successful in identifying their large-scale structures, e.g. for node clustering [7]. Several papers have therefore used tools from random graph theory and graphons to study the behavior of GNNs *in the large-graph limit*. In [22, 38], GNNs are shown to converge to limiting “continuous” architectures (coined c-GNNs in [22]). A few works have studied the discriminative power of c-GNNs on graphons [30], however, analogously to how the WL test will fail on regular graphs, c-GNNs are severely limited on graph models with almost-constant degree function (Fig. 1), and the question is still largely open.

Contribution and outline. In this paper, we study the convergence of SGNNs on large random graphs towards “c-SGNNs”, and analyze the approximation power of c-GNNs and c-SGNNs. After some preliminary results in Sec. 2, we study the convergence of SGNNs in Sec 3. In Sec. 4 and 5, we show that c-SGNNs are strictly more powerful than c-GNNs in both permutation-invariant and equivariant case. We then prove the universality of c-SGNNs on several popular models of random graphs, including Stochastic Block Models (SBMs) and radial kernels. For instance, we show in Sec. 5.3 that c-SGNNs can perfectly identify the communities of most SBMs, including some for which any c-GNN provably fails (Fig. 1).

Related work. The approximation power of GNNs (on finite graphs) has been a topic of great interest in recent years, and we do not attempt to make an exhaustive list of all results here. The landmark paper [45] showed that permutation-invariant GNNs were at best as powerful as the WL test, and later models [31, 8, 16] were constructed to be as powerful as higher-order WL tests, using for instance higher-order tensors [31]. The link between the graph isomorphism problem and function approximation was made in [8, 23] and extended in [3]. Some architectures were proven to be universal [32, 23] but involve tensors of unbounded order. When graphs are equipped with unique node identifiers, the approximation power of GNNs can be significantly improved if relaxing permutation-invariance/equivariance [28, 27, 10], or assuming that it holds only in expectation with respect to additional sources of randomness [39, 1]. On the contrary, as we will see in Sec. 2.3, the SGNN architecture proposed in [41] is inspired by the same ideas but still satisfies exact permutation-invariance/equivariance. As mentioned above, SGNNs are more powerful than the WL test [41], and even universal on graphs with bounded degrees when allowing powerful layers. To our knowledge, the approximation power of SGNNs as they are implemented in practice is still open. We treat of their continuous limit in this paper.

Fewer results can be found in the large-graph limit. Beyond the graph-isomorphism paradigm, authors have studied the capacity of GNNs to count graph substructures [6, 13] or identify various graph properties [15, 46]. Several works have studied the large-graphs limit of GNNs [38, 30, 26, 22], assuming random graph or graphon models. The degree function has been proven to be a crucial element for the discriminative power of c-GNNs, and they will not be able to distinguish graphons with close degree functions [30]. Here we show how c-SGNNs allow to overcome these limitations, and provide the first universality theorems for (S)GNNs in the continuous limit. Finally, we note that universal architectures exist on *measures* [11, 49], which can be seen as limits of *point clouds*, but, similar to the discrete case [48, 32, 23], the graph case may be significantly harder to study.




Figure 1: Illustration of Prop. 7 (Sec. 5.3). On an SBM with constant degree function, a GNN (top) might overfit the training set, but converges to a constant c-GNN. On the contrary, there exists a c-SGNN (bottom) that perfectly separates the communities. Details can be found in App. E.

2 Preliminaries

In this section, we group notations on random graphs and GNNs, present the convergence result of [22] slightly adapted to our context, and introduce the SGNNs of [41]. An undirected graph G with n nodes is represented by a symmetric adjacency matrix $A \in \mathbb{R}^{n \times n}$. It is isomorphic to all graphs with adjacency matrices $\sigma A \sigma^\top$, for any permutation matrix $\sigma \in \{0, 1\}^{n \times n}$. Matrix rows (resp. columns) are denoted by $A_{i,:}$ (resp. $A_{:,i}$). The norm $\|\cdot\|$ is the operator norm for matrices and Euclidean norm for vectors. For a compact metric space \mathcal{X} , we denote by $\mathcal{C}(\mathcal{X}, \mathbb{R})$ the space of bounded continuous functions $\mathcal{X} \rightarrow \mathbb{R}$, equipped with $\|f\|_\infty = \sup_x |f(x)|$. For a multivariate function $f \in \mathcal{C}(\mathcal{X}, \mathbb{R}^d)$, f_i denotes its i th coordinate, and $\|f\|_\infty = (\sum_i \|f_i\|_\infty^2)^{1/2}$.

2.1 Random graph model

We consider ‘‘latent position’’ random graphs, which include several popular models such as SBMs or graphons [29]. The nodes are associated with unobserved latent random variables x_i drawn *i.i.d.* For the edges, we will examine two cases: an ‘‘ideal’’ one with deterministic weighted edges, and a more ‘‘realistic’’ one where the edges are randomly drawn independently. In the literature, the latter is often considered as an idealized framework to study the behavior of various algorithms [42, 37].

Let $(\mathcal{X}, m_{\mathcal{X}})$ be a compact metric space that is not a singleton, and assume that the ε -covering numbers¹ of \mathcal{X} scale as $\mathcal{O}(\varepsilon^{-d_{\mathcal{X}}})$ for some dimension $d_{\mathcal{X}}$. We denote by \mathcal{W} the set of symmetric bivariate functions in $\mathcal{C}(\mathcal{X} \times \mathcal{X}, [0, 1])$ that are L_W -Lipschitz in each variable, and by \mathcal{P} the set of probability distributions over \mathcal{X} equipped with the total variation norm $\|\cdot\|_{\text{TV}}$. A graph with n nodes is generated according to a random graph model $(W, P) \in \mathcal{W} \times \mathcal{P}$ as follows:

$$x_1, \dots, x_n \stackrel{i.i.d.}{\sim} P, \quad a_{ij} = \begin{cases} W(x_i, x_j) & \text{deterministic edges case} \\ \alpha_n^{-1} \text{Bernoulli}(\alpha_n W(x_i, x_j)) & \text{random edges case} \end{cases}$$

where α_n is the *sparsity level* of the graph in the random edges case, which we assume to be known for simplicity. When $\alpha_n \sim 1$, the graph is said to be *dense*, when $\alpha_n \sim \frac{1}{n}$ the graph is *sparse*, and when $\alpha \sim \frac{\log n}{n}$ the graph is *relatively sparse*. Note that we have normalized the Bernoulli edges by $1/\alpha_n$ such that $\mathbb{E}a_{ij} = W(x_i, x_j)$ conditionally on x_i, x_j . When \mathcal{X} is finite, the model is called an SBM (see Sec. 4.3).

Like finite graphs, random graph models can be isomorphic to one another [22, 29]. In this paper, similar to [22] we consider that, for any bijection $\varphi : \mathcal{X} \rightarrow \mathcal{X}$, the model (W, P) is isomorphic to $(W_\varphi, \varphi_\# P)$, where $W_\varphi(x, y) = W(\varphi(x), \varphi(y))$ and $\varphi_\# P$ is the pushforward of P (the distribution of $f(X)$). Indeed, it is easy to see that both produce exactly the same distribution over graphs.

2.2 Graph Neural Networks

Following [22, 12], in this paper we consider the so-called ‘‘spectral’’ version of GNNs, which include several message-passing models for certain aggregation functions. We consider polynomial filters h defined as $h(A) = \sum_k \beta_k A^k$ for a matrix or operator A . In practice, the order of the filters is always finite, but our results are valid for infinite-order filters (assuming that the sum always converges for simplicity). We consider any activation function $\rho : \mathbb{R} \rightarrow \mathbb{R}$ which satisfies $\rho(0) = 0$ and $|\rho(x) - \rho(y)| \leq |x - y|$ for which the universality theorem of MLPs applies [35], e.g. ReLU.

Spectral GNNs are defined by successive filtering of a graph signal. Given an input signal $Z^{(0)} \in \mathbb{R}^{n \times d_0}$, at each layer $\ell = 0, \dots, M - 1$:

$$Z_{:,j}^{(\ell+1)} = \rho \left(\sum_{i=1}^{d_\ell} h_{ij}^{(\ell)} \left(\frac{1}{n} A \right) Z_{:,i}^{(\ell)} + b_j^{(\ell)} 1_n \right) \in \mathbb{R}^n \quad j = 1, \dots, d_{\ell+1}, \quad (1)$$

where $h_{ij}^{(\ell)}(A) = \sum_k \beta_{ijk}^{(\ell)} A^k$ are trainable graph filters, $b_j^{(\ell)} \in \mathbb{R}$ are trainable additive biases and ρ is applied pointwise. We note the normalization A/n that will be necessary for convergence. GNNs exist in two main versions: so-called ‘‘permutation-invariant’’ GNNs $\bar{\Phi}$ output a single vector for the entire graph, while ‘‘permutation-equivariant’’ GNNs Φ output a graph signal:

$$\Phi_A(Z^{(0)}) = g(Z^{(M)}) \in \mathbb{R}^{n \times d_{out}} \quad \bar{\Phi}_A(Z^{(0)}) = g\left(\frac{1}{n} \sum_{i=1}^n Z_{i,:}^{(M)}\right) \in \mathbb{R}^{d_{out}}, \quad (2)$$

¹the number of balls of radius ε required to cover \mathcal{X}

where $g : \mathbb{R}^{d_M} \rightarrow \mathbb{R}^{d_{out}}$ is an MLP applied row-wise in the equivariant case. By construction, such GNNs are indeed equivariant or invariant to node permutation: for all permutation matrices σ , we have $\Phi_{\sigma A \sigma^\top}(\sigma Z^{(0)}) = \sigma \Phi_A(Z^{(0)})$ and $\bar{\Phi}_{\sigma A \sigma^\top}(\sigma Z^{(0)}) = \bar{\Phi}_A(Z^{(0)})$.

Continuous GNNs. In [22], the authors show that, in the large random graphs limit, GNNs converge to the following ‘‘continuous’’ models (coined c-GNNs), which propagate **functions over the latent space** $f^{(\ell)} \in \mathcal{C}(\mathcal{X}, \mathbb{R}^{d_\ell})$ instead of graph signals. Given an input function $f^{(0)} \in \mathcal{C}(\mathcal{X}, \mathbb{R}^{d_0})$:

$$f_j^{(\ell+1)} = \rho \left(\sum_{i=1}^{d_\ell} h_{ij}^{(\ell)}(T_{W,P}) f_i^{(\ell)} + b_j^{(\ell)} \right),$$

where $T_{W,P}$ is the operator $T_{W,P} f = \int W(\cdot, x) f(x) dP(x)$. Then similar to (2) the permutation-equivariant/invariant versions of c-GNNs are defined as:

$$\Phi_{W,P}(f^{(0)}) = g \circ f^{(M)}, \quad \bar{\Phi}_{W,P}(f^{(0)}) = g \left(\int f^{(M)}(x) dP(x) \right).$$

Remark that $\Phi_{W,P}(f^{(0)})$ is itself a *function* in $\mathcal{C}(\mathcal{X}, \mathbb{R}^{d_{out}})$ while $\bar{\Phi}_{W,P}(f^{(0)}) \in \mathbb{R}^{d_{out}}$ is a vector. Moreover, by virtue of the polynomial filters, the four architectures $\Phi_A, \bar{\Phi}_A, \Phi_{W,P}, \bar{\Phi}_{W,P}$ have the exact same set of parameters. Like in the discrete case, one can prove that c-GNN are equivariant or invariant to isomorphisms of random graph models [22]: for all bijection $\varphi : \mathcal{X} \rightarrow \mathcal{X}$, we have $\Phi_{W_\varphi, \varphi_\#^{-1} P}(f \circ \varphi) = \Phi_{W,P}(f) \circ \varphi$ and $\bar{\Phi}_{W_\varphi, \varphi_\#^{-1} P}(f \circ \varphi) = \bar{\Phi}_{W,P}(f)$.

Let us now turn to convergence of GNNs to c-GNNs. The following result is adapted² from [22]. While the outputs of permutation-invariant GNNs and c-GNNs are vectors in $\mathbb{R}^{d_{out}}$ that can be directly compared, the output graph signal of a permutation-equivariant GNN is compared with a *sampling* of the output function of the corresponding c-GNN: for a graph signal $Z \in \mathbb{R}^{n \times d}$, a function $f : \mathcal{X} \rightarrow \mathbb{R}^d$ and $X = \{x_i\}_{i=1}^n$, we define the (square root of the) mean-square error as $\text{MSE}_X(Z, f) = \left(\frac{1}{n} \sum_{i=1}^n \|Z_i - f(x_i)\|_2^2 \right)^{1/2}$. The proof of Theorem 1 with all multiplicative constants is given in App. A.1. Recall that d_X is the ‘‘dimension’’ of \mathcal{X} .

Theorem 1. *Assume G is drawn from (W, P) and has latent variables X . Fix $\rho, \nu > 0$.*

- **In the deterministic edges case:** with probability $1 - \rho$, for all $Z^{(0)}$:

$$\text{MSE}_X(\Phi_A(Z^{(0)}), \Phi_{W,P}(f^{(0)})) \leq C \cdot \text{MSE}_X(Z^{(0)}, f^{(0)}) + R_1(n) \quad (3)$$

for some constant C and $R_1(n) = \mathcal{O} \left(\sqrt{(d_X + \log(1/\rho))/n} \right)$.

- **In the random edges case:** assume that the sparsity level is $\alpha_n \gtrsim n^{-1} \log n$. There is a constant C_ν such that, with probability $1 - \rho - n^{-\nu}$, for all $Z^{(0)}$:

$$\text{MSE}_X(\Phi_A(Z^{(0)}), \Phi_{W,P}(f^{(0)})) \leq C \cdot \text{MSE}_X(Z^{(0)}, f^{(0)}) + R_1(n) + R_2(n) \quad (4)$$

where $R_2(n) = \mathcal{O} \left(C_\nu / \sqrt{\alpha_n n} \right)$.

- **In the permutation-invariant case:** The exact same results hold for $\|\bar{\Phi}_A(Z^{(0)}) - \bar{\Phi}_{W,P}(f^{(0)})\|$ instead of the MSE on the left-hand-side, with an added error term $R_3(n) = \mathcal{O} \left(\sqrt{\log(1/\rho)/n} \right)$.

By the theorem above, a GNN converges to its continuous counterpart if $Z^{(0)}$ is (close to) a sampling of a function $f^{(0)}$ at the latent variables. This is directly assumed in [22]. In the present paper however, we do not suppose that input node features are available. While several strategies have been proposed in the literature, a popular baseline is to simply take constant input $Z^{(0)} = 1_n$ (which, by a multiplication by A on the first layer, is also equivalent to inputting the degrees as in [7] for instance). In this case, there is convergence to a c-GNN with $f^{(0)} = 1$. For simplicity in the rest of the paper we drop the notation ‘‘(1)’’ and write $\Phi_A = \Phi_A(1_n)$, $\Phi_{W,P} = \Phi_{W,P}(1)$, and so on. It is known that GNNs are limited on regular graphs: for permutation-invariant GNNs, the WL test cannot distinguish regular graphs of the same order, and for permutation-equivariant GNNs, $\Phi_A(1_n)$ is *constant over the nodes*. Similarly for c-GNN, the degree function $\int W(\cdot, x) dP(x)$ is key in the discriminative power of c-GNNs [30], and if it is constant, then $\Phi_{W,P}(1)$ is a constant function (see Fig. 1).

²The authors in [22] proved this for the normalized Laplacian, which allows bypassing the knowledge of α_n . Here we use the adjacency matrix, for our later results on approximation power.

2.3 SGNN: GNN with unique node identifiers

To remedy the absence of input node features, in [41] the authors propose an architecture with *unique node identifiers* while still respecting permutation invariance/equivariance. More precisely, they first choose an *arbitrary* ordering of the nodes $q = 1, \dots, n$, apply a GNN to each *one-hot vector* $e_q = [0, \dots, 1, \dots, 0]$, and restore equivariance to permutation by a final pooling. For later purpose of convergence, we generalize this strategy to any collection $E_q(A) \in \mathbb{R}^{n \times d_0}$ that satisfies:

$$E_q(\sigma A \sigma^\top) = \sigma E_{\sigma^{-1}(q)}(A). \quad (5)$$

where by an abuse of notation $\sigma(q)$ designates the permutation function applied to index q . For instance, it can be any filtering of one-hot vector $E_q(A) = h(A)e_q$. After the GNN Φ_A , a pooling is applied to restore equivariance, then a second GNN Φ'_A that is either invariant or equivariant:

$$\Psi_A = \Phi'_A \left(\frac{1}{n} \sum_q \Phi_A(E_q(A)) \right) \in \mathbb{R}^{n \times d'_{out}}, \quad \bar{\Psi}_A = \bar{\Phi}'_A \left(\frac{1}{n} \sum_q \Phi_A(E_q(A)) \right) \in \mathbb{R}^{d'_{out}} \quad (6)$$

In [41], these architectures, called SMPs, are interpreted as doing message-passing over matrices, and can use more general pooling and aggregation functions. In a sense, what we call SGNN are “spectral” versions of SMP, but are essentially the same idea. It is not difficult to see that SGNNs satisfy: $\Psi_{\sigma A \sigma^\top} = \sigma \Psi_A$ and $\bar{\Psi}_{\sigma A \sigma^\top} = \bar{\Psi}_A$.

3 Convergence of SGNNs on large random graphs

In this section, we extend Theorem 1 to SGNNs. To define continuous SGNNs, we consider a *bivariate* input function $\eta_{W,P} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{d_0}$ such that: the mapping $(W, P) \mapsto \eta_{W,P}$ is continuous³, for any (W, P) , $\eta_{W,P}$ is C_η -bounded and L_η -Lipschitz in each variable, and similar to (5) it respects the following:

$$\eta_{W_\varphi, \varphi_\#^{-1}P}(x, y) = \eta_{W,P}(\varphi(x), \varphi(y)), \quad (7)$$

for all bijections $\varphi : \mathcal{X} \rightarrow \mathcal{X}$. For instance, $\eta_{W,P} = W$ or any filter $\eta_{W,P}(x, y) = [h(T_{W,P})W(\cdot, y)](x)$ satisfy these conditions. A c-SGNN is then defined as:

$$\Psi_{W,P} = \Phi'_{W,P} \left(\int \Phi_{W,P}(\eta_{W,P}(\cdot, x)) dP(x) \right), \quad (8)$$

for the equivariant case, or similarly $\bar{\Psi}_{W,P} = \bar{\Phi}'_{W,P} \left(\int \Phi_{W,P}(\eta_{W,P}(\cdot, x)) dP(x) \right)$ for the invariant case. Again, using the properties of c-GNNs, it is easy to see that c-SGNNs respect random graphs isomorphism: $\Psi_{W_\varphi, \varphi_\#^{-1}P} = \Psi_{W,P} \circ \varphi$ and $\bar{\Psi}_{W_\varphi, \varphi_\#^{-1}P} = \bar{\Psi}_{W,P}$. We are now ready to extend Theorem 1. The proof and multiplicative constants are in App. A.2.

Theorem 2. *Assume G is drawn from (W, P) with latent variables X . Fix $\rho, \nu > 0$.*

- **In the deterministic edges case:** with probability $1 - \rho$:

$$\text{MSE}_X(\Psi_A, \Psi_{W,P}) \leq C' \cdot \sup_q \text{MSE}_X(E_q(A), \eta(\cdot, x_q)) + R'_1(n) \quad (9)$$

for some constant C' and $R'_1(n) = \mathcal{O} \left(\sqrt{(d_{\mathcal{X}} + \log(1/\rho))/n} \right)$.

- **In the random edges case:** assume that the sparsity level is $\alpha_n \gtrsim n^{-1} \log n$. There is a constant C_ν such that, with probability $1 - \rho - n^{-\nu}$:

$$\text{MSE}_X(\Psi_A, \Psi_{W,P}) \leq C' \cdot \sup_q \text{MSE}_X(E_q(A), \eta(\cdot, x_q)) + R'_1(n) + R'_2(n) \quad (10)$$

where $R'_2(n) = \mathcal{O} \left(C_\nu / \sqrt{\alpha_n n} \right)$.

- **In the permutation-invariant case:** The exact same results hold for $\|\bar{\Psi}_A - \bar{\Psi}_{W,P}\|$ instead of the MSE, with an added error term $R'_3(n) = \mathcal{O} \left(\sqrt{\log(1/\rho)/n} \right)$.

Hence, we obtain convergence when the input signal $E_q(A)$ is close to being a sampling of a function $\eta_{W,P}$ at x_q . The choice of input signal/function is therefore crucial. Let us examine a few strategies.

³for the norm $\|\cdot\|_\infty + \|\cdot\|_{\text{TV}}$ on $\mathcal{W} \times \mathcal{P}$

One-hot vectors. If one chooses one-hot vectors $E_q = e_q$ as in [41], then we can see that the SGNN converges to the continuous architecture with input $\eta_{W,P} = 0$, since $\text{MSE}_X(e_q, 0) \rightarrow 0$. Since $\Psi_{W,P}$ is nothing more than a traditional c-GNN $\Phi_{W,P}$ with constant input in that case, this is not a satisfying choice in terms of approximation power.

One-hop filtering. If we choose to “filter” e_q once and take $E_q(A) = Ae_q$, the natural continuous equivalent is $\eta_{W,P} = W$. Such a strategy only works for deterministic edges: indeed,

$$\text{MSE}_X(Ae_q, W(\cdot, x_q)) = (n^{-1} \sum_{i=1}^n (a_{iq} - W(x_i, x_q))^2)^{1/2}$$

$$\begin{cases} = 0 & \text{with deterministic edges} \\ \approx \sqrt{n^{-1} \sum_i \text{Var}(a_{iq})} \sim \alpha^{-1} & \text{with random edges, w.h.p.} \end{cases}$$

where the last line comes from a simple application of Hoeffding’s inequality. Hence, in the case of random edges, the MSE does not vanish and typically diverges for non-dense graphs (Fig. 2).

Two-hop filtering. We can therefore choose to filter *twice* and consider $\eta_{W,P}(x, y) = T_{W,P}[W(\cdot, y)](x)$, that is, $E_q(A) = A^2 e_q/n$. We have the following result.

Proposition 1. *In the random edges case, with probability $1 - \rho$, we have for all q :*

$$\text{MSE}_X \left(\frac{A^2 e_q}{n}, T_{W,P}[W(\cdot, x_q)](\cdot) \right) \lesssim \mathcal{O} \left(\frac{\sqrt{\log(n/\rho)}}{\alpha_n \sqrt{n}} \right).$$

In the deterministic edges case, the rate is $\mathcal{O}(1/\sqrt{n})$.

Convergence is guaranteed when $\sqrt{\log n}/(\alpha_n \sqrt{n}) \rightarrow 0$, which is stronger than relative sparsity. The difference between one-hop and two-hop filtering is illustrated in Fig. 2. Other strategies may also be examined, and we leave this for future work. For instance, based on Theorem 5 in App. D, a strategy based on the eigenvectors of A could lead to convergence even in the relatively sparse case. In the rest of the paper, we explicitly write when our results are valid for one- or two-hop filtering.




Figure 2: Difference between the outputs of some $\bar{\Psi}_A$ in the random edges case with $\alpha_n \sim n^{-1/3}$ and in the deterministic edges case, which converges to $\bar{\Psi}_{W,P}$. Convergence is observed for two-hop filtering only. The theoretical rate of Prop. 1 is slightly pessimistic. Details can be found in App. E.

4 Approximation power: permutation-invariant case

In this section, we study the approximation power of continuous architectures in the permutation-invariant case. We seek to characterize the functions $\mathcal{W} \times \mathcal{P} \rightarrow \mathbb{R}$ that can be well-approximated by a c-GNN or c-SGNN. We derive a generic criterion for universality arising from the Stone-Weierstrass theorem, before proving that c-SGNNs are indeed strictly more powerful than c-GNNs. Finally, we give several examples of models for which c-SGNNs are universal.

4.1 Generic result with Stone-Weierstrass theorem

A classical tool to prove universality of neural nets in the literature is the Stone-Weierstrass theorem [20, 23], which states that an algebra of functions that separates points on a compact space is dense in the set of continuous functions. Although NNs are typically not an algebra since they are not closed by multiplication, one of the original proofs of universality for MLPs solves this by a clever trick [20] which we recall in Lemma 1 in App. D. A direct application results in the following proposition.

Proposition 2. *Let $\mathcal{M} \subset \mathcal{W} \times \mathcal{P}$ be a compact subset of $\mathcal{W} \times \mathcal{P}$. c-SGNNs (resp. c-GNNs) are dense in $\mathcal{C}(\mathcal{M}, \mathbb{R})$ if and only if: for all $(W, P) \neq (W', P') \in \mathcal{M}$, there is a c-SGNN (resp. a c-GNN) such that $\bar{\Psi}_{W,P} \neq \bar{\Psi}_{W',P'}$ (resp. $\bar{\Phi}_{W,P} \neq \bar{\Phi}_{W',P'}$).*

Note that, by construction of the permutation-invariant c-(S)GNNs, universality is only possible when \mathcal{M} does not contain two isomorphic versions of the same random graph model. Equivalently, \mathcal{M} may be a larger set *quotiented* by random graph isomorphism.

4.2 c-SGNNs are more powerful than c-GNNs

SGNNs were proven to be strictly more powerful than the WL test, and therefore than GNNs, in [41]. In the theorem below, we check that this strict inclusion holds for their continuous limits.

Theorem 3. *The set of functions of the form $(W, P) \rightarrow \bar{\Phi}_{W,P}$ is **strictly** included in the set of functions $(W, P) \rightarrow \bar{\Psi}_{W,P}$, for both one- and two-hop input filtering.*

This theorem is proven by constructing two models that are distinguished by a c-SGNN but not by any c-GNN. We note that there might be subsets $\mathcal{M} \subset \mathcal{W} \times \mathcal{P}$ that do not contain such models, and therefore on which c-GNNs and c-SGNNs have the same approximation power.

4.3 Examples

While a generic universality theorem on random graphs seems to be out-of-reach for the moment, we examine several interesting examples, and pave the way for future extensions. We focus on c-SGNNs, but, sometimes, may not conclude on the power of c-GNNs: we could not prove that they are universal, but were not able to find a counter-example either. For simplicity, and since our purpose here is mainly illustrative, we mostly focus on one-hop input filtering.

Stochastic Block Models. SBMs [19] are classical models to emulate graphs with communities. In our settings, SBMs with K communities can be obtained with a finite latent space $|\mathcal{X}| = K$, typically $\mathcal{X} = \{1, \dots, K\}$. The kernel $W(k, k') = W_{kk'}$ can be represented as a matrix $W \in S_K$, where S_K is the set of symmetric matrices in $[0, 1]^{K \times K}$, and the distribution $P(k) = P_k$ as a vector $P \in \Delta^{K-1}$, where $\Delta^{K-1} = \{P \in [0, 1]^K, \sum_k P_k = 1\}$ is the $(K-1)$ -dimensional simplex.

In the following proposition, we fix P and examine universality with respect to W . In this case, continuous GNNs are actually quite similar to discrete ones on matrices W , except that the probability vector P also intervenes in the computation. While GNNs on finite graphs can only be universal when using high-order tensors [32, 23] due to invariance to graph-isomorphism, here P can help to disambiguate this constraint. We will say that $P \in \Delta^{K-1}$ is **incoherent** if: for signs $s \in \{-1, 0, 1\}^k$, having $\sum_{k=1}^K s_k P_k = 0$ implies $s = 0$. That is, no probability is an exact sum or difference of the others. We note that a vector drawn uniformly on Δ^{K-1} is incoherent with probability 1. For incoherent probability vectors, we can show universality of c-SGNNs. Moreover, this is actually a case where we *can* prove that c-GNNs are, in turn, *not* universal.

Proposition 3. *For one-hop input filtering, if P is incoherent the space of functions $W \rightarrow \bar{\Psi}_{W,P}$ is dense in $\mathcal{C}(S_K, \mathbb{R})$. Moreover, there exists P incoherent and $W \neq W'$ such that, for any c-GNN, $\bar{\Phi}_{W,P} = \bar{\Phi}_{W',P}$.*

Additive kernel. Let us now fix W and examine universality with respect to P . A classical theorem on symmetric continuous functions [48] states that any W can be arbitrarily well approximated as $W(x, y) \approx u(v(x) + v(y))$ for some functions u, v . Inspired by this result, a kernel will be said to be **additive** if it can (exactly) be written as $W(x, y) = u(v(x) + v(y))$, and **injectively additive** if both u, v are *continuous and injective*. We prove universality in the unidimensional case below.

Proposition 4. *Let $\mathcal{X} \subset \mathbb{R}$, and $\tilde{\mathcal{P}}$ be any compact subset of \mathcal{P} . Assume W is injectively additive with $\text{Im}(v) \subset \mathbb{R}$. For one-hop filtering, the space of functions $P \rightarrow \bar{\Psi}_{W,P}$ is dense in $\mathcal{C}(\tilde{\mathcal{P}}, \mathbb{R})$.*

It is easy to see that injectively additive kernels include all SBMs for which $W_{ij} \neq W_{i'j'}$ when $j \neq j'$, so Prop. 4 completes Prop. 3. However, unlike additive kernels, injectively additive kernels are *a priori* **not** universal approximators of symmetric continuous functions: this result [48] is only valid when $\text{Im}(v)$ can be multidimensional. But it is known for instance that there is no continuous injective map from $[0, 1]^2$ to $[0, 1]$, so if $\text{Im}(v) = [0, 1]^2$, u cannot be both continuous and injective.

Radial kernel. We conclude this section with an important class of kernels, *radial* kernels $W(x, y) = w(\|x - y\|)$ for some function $w : \mathbb{R}_+ \rightarrow [0, 1]$. They include the popular Gaussian kernel and so-called ε -graphs. Below, we give an example in one dimension, for which c-SGNNs are universal on symmetric distributions. The case of non-symmetric distributions seems more involved and we leave it for future investigations.

Proposition 5. Assume that $\mathcal{X} = [-1, 1]$ and $W(x, y) = w(|x - y|)$ where w is continuous and injective. Let $\tilde{\mathcal{P}} \subset \mathcal{P}$ be any compact set of **symmetric** distributions. For one-hop input filtering, the space of functions $P \rightarrow \bar{\Psi}_{W,P}$ is dense in $\mathcal{C}(\tilde{\mathcal{P}}, \mathbb{R})$.

5 Approximation power: permutation-equivariant case

In the equivariant case, recall that the outputs of c-(S)GNNs are *functions* on \mathcal{X} . The “traditional” notion of universality is to evaluate the approximation power of mappings $(W, P) \rightarrow \mathcal{F}$, where \mathcal{F} is some space of equivariant functions, as is done for the discrete case in [23, 3]. However, a potentially simpler and more relevant notion here is to *fix* (W, P) , and directly examine the properties of the space of functions $\mathcal{X} \rightarrow \mathbb{R}$ represented by c-(S)GNNs, that is, the space of functions $\{\Psi_{W,P} : \mathcal{X} \rightarrow \mathbb{R}\}$ for all possible c-SGNNs $\Psi_{W,P}$ (and similar for $\Phi_{W,P}$). Indeed, this directly answers such questions as: given an SBM, does there exist a c-GNN that can label the communities (Fig. 1)? Or: given the structure of a mesh, what functions can be computed on it, e.g. for segmentation?

5.1 Generic result with Stone-Weierstrass theorem

As in the invariant case, the Stone-Weierstrass theorem yields a generic separation condition.

Proposition 6. Let (W, P) be fixed. Then c-SGNNs (resp. c-GNNs) are dense in $\mathcal{C}(\mathcal{X}, \mathbb{R})$ iff: for all $x \neq x' \in \mathcal{X}$, there is a c-SGNN (resp. a c-GNN) such that $\Psi_{W,P}(x) \neq \Psi_{W,P}(x')$ (resp. $\Phi_{W,P}(x) \neq \Phi_{W,P}(x')$).

If (W, P) are such that c-(S)GNNs satisfy some symmetry or invariance (see for instance Prop. 9), then the space \mathcal{X} can be quotiented to obtain universality among functions satisfying these constraints.

5.2 c-SGNNs are more powerful than c-GNNs

Using a proof similar to Theorem 3, we can then prove that c-SGNNs are indeed strictly more powerful than c-GNNs for some (W, P) .

Theorem 4. For both one- and two-hop input filtering, the following holds. For any (W, P) , the set of functions of the form $\Phi_{W,P}$ is included in the set of functions $\Psi_{W,P}$, and there exist (W, P) such that the inclusion is **strict**.

Again, we note that, for some random graph models (W, P) , c-SGNNs and c-GNNs might have the same approximation power.

5.3 Examples

We treat the same examples as before, with the addition of two-hop filtering for SBMs and radial kernels on the d -dimensional sphere.

SBM. Universality in the SBM case corresponds to being able to distinguish communities, that is, $\Psi_{W,P} : \mathcal{X} \rightarrow \mathbb{R}$ returns a different value for each element of the latent space \mathcal{X} . In the following result, we assume that W is invertible, and prove the result for *both* one- and two-hop filtering, meaning that c-SGNNs can indeed distinguish communities of SBMs with random edges under some mild conditions. On the other hand, c-GNNs may fail on such models.

Proposition 7. Let P be incoherent, and W be invertible. For **both** one- and two-hop input filtering, c-SGNNs are dense in $\mathcal{C}(\mathcal{X}, \mathbb{R})$. Moreover, there exist (W, P) satisfying the conditions above such that any c-GNN $\Phi_{W,P}$ is a constant function.

This proposition is illustrated in Fig. 1: on an SBM with constant degree function, any GNN converges to a c-GNN with constant output, while a SGNN with two-hop input filtering can be close to a c-SGNN that perfectly separates the communities.

Additive kernel. As in the invariant case, injectively additive kernels lead to universality.

Proposition 8. Assume W is injectively additive, fix any P . Then, for one-hop input filtering, c-SGNNs are dense in $\mathcal{C}(\mathcal{X}, \mathbb{R})$.




Figure 3: Illustration of Prop. 9 on Gaussian kernel, but with random edges and two-hop input filtering. The x-axis is the latent variables in $\mathcal{X} = [-1, 1]$. The y-axis is the output of a SGNN trained to approximate some function f (red curve). On the left, both distribution P and f are symmetric, c-SGNNs are universal in that case. In the center, P is symmetric but not f , and the training expectedly fails since the limit c-SGNN is symmetric. On the right, P and f are non-symmetric, and universality holds again. Details can be found in App. E.

Radial kernel. Unlike the invariant case (Prop. 5) which was limited to symmetric distributions, we treat both symmetric and non-symmetric case: when P is symmetric, then so is $\Psi_{W,P}$ (by permutation-equivariance), and we have universality *among symmetric functions*. When P is non-symmetric, we have universality among *all* functions. See Fig. 3 for an illustration.

Proposition 9. Consider $\mathcal{X} = [-1, 1]$, $W(x, y) = w(|x - y|)$ a radial kernel with an invertible, continuous w , and P with a piecewise continuous density such that $\mathbb{E}_P X = 0$. Then, for one-hop input filtering: if P is symmetric, c-SGNNs are dense in the space of symmetric functions in $\mathcal{C}(\mathcal{X}, \mathbb{R})$, and if P is not symmetric, c-SGNNs are dense in $\mathcal{C}(\mathcal{X}, \mathbb{R})$.

Finally, we look at radial kernels on the d -dimensional sphere $\mathcal{X} = \mathbb{S}^{d-1} = \{x \in \mathbb{R}^d \mid \|x\| = 1\}$, an important example sometimes referred to as random geometric graphs [34, 2], or dot-product kernels [33], which are for instance popular in social networks analysis [47]. Indeed, in this case, the kernel only depends on the dot product $W(x, y) = w(x^\top y)$. Denoting by $d\tau$ the uniform measure on \mathbb{S}^{d-1} , it is known [14, 9] that functions in $L^2(d\tau)$ can be uniquely decomposed as $f(x) = \sum_{k \geq 0} f_k(x) = \sum_{k \geq 0} \sum_{j=1}^{N(d,k)} a_{k,j} Y_{k,j}(x)$ where $Y_{k,j}$ are *spherical harmonics*, that is, homogeneous harmonic polynomials of degree k which form an orthonormal basis of $L^2(d\tau)$. We will say that such a function is **injectively decomposed** if the mapping $x \rightarrow [f_k(x)]_{k \geq 0}$ from \mathbb{S}^{d-1} to $\ell_2(\mathbb{R})$ is injective. Note that generically this is verified if f_k is non-zero for more than $d - 1$ distinct values of $k > 0$, as this corresponds to solving an over-determined system of polynomial equations, but there may be degenerate situations where this is not enough. The proof of the following proposition is based on the well-known Legendre/Gegenbauer polynomial decomposition of spherical harmonics [9] (see App. C.5).

Proposition 10. Assume that $\mathcal{X} = \mathbb{S}^{d-1}$, that $W(x, y) = w(x^\top y)$ with continuous invertible $w : [-1, 1] \rightarrow [0, 1]$, and that $P = f d\tau$ has a density f which is injectively decomposed. Then for one-hop input filtering c-SGNNs are dense in $\mathcal{C}(\mathcal{X}, \mathbb{R})$.

6 Conclusion and outlooks

It is known that permutation-invariant GNNs fail to distinguish regular graphs of the same order, and permutation-equivariant GNNs return constant output on regular graphs. Similarly, their continuous counterparts suffer from the same flaw on random graph with constant or almost-constant degree function [30]. However, we showed that the recently proposed SGNNs converge to continuous architectures which, like in the discrete world, are strictly more powerful than c-GNNs. Moreover, we proved that both permutation-invariant and permutation-equivariant c-SGNNs are universal on many random graph models of interest, including a large class of SBMs and random geometric graphs.

We believe that our work opens many possibilities for future investigations. We examined very simple strategies for choosing the inputs $E_q(A)$ of the SGNN, but more complex, spectral-based choices could exhibit better convergence properties and approximation power. We showed universality in specific random graph models of interest, but deriving a more generic criterion is still an open question. More directly, most of our examples illustrate the one-dimensional case $\mathcal{X} \subset \mathbb{R}$, and a generalization to multidimensional latent spaces would be an important step forward. Besides SGNNs, architectures that include high-order tensors [31] (sometimes called FGNN [3]) are known

to be more powerful than the WL test. Conditions for their convergence on large graphs are still open, in particular since high-order tensors lead to high-order operators that may be difficult to manipulate. Finally, we remark that directly estimating the latent variables x_i is a classical task in statistics, for which conditions of success have been derived for various approaches, e.g. for Spectral Clustering [25]. Comparing them with (S)GNNs is an important path for future work.

Acknowledgement

This work was partly supported by ANR GraVa ANR-18-CE40-0005.

References

- [1] Ralph Abboud, İsmail İlkan Ceylan, Martin Grohe, and Thomas Lukasiewicz. The Surprising Power of Graph Neural Networks with Random Node Initialization. In *IJCAI International Joint Conference on Artificial Intelligence*, pages 2112–2118, 2021.
- [2] Ernesto Araya and Yohann de Castro. Latent distance estimation for random geometric graphs. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [3] Waiss Azizian and Marc Lelarge. Characterizing the expressive power of invariant and equivariant graph neural networks. In *International Conference on Learning Representations (ICLR)*, 2021.
- [4] László Babai. Graph isomorphism in quasipolynomial time. *Proceedings of the Annual ACM Symposium on Theory of Computing*, 2016.
- [5] Michael M. Bronstein, Joan Bruna, Yann Lecun, Arthur Szlam, and Pierre Vandergheynst. Geometric Deep Learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [6] Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. Can graph neural networks count substructures? In *Advances in Neural Information and Processing Systems (NeurIPS)*, 2020.
- [7] Zhengdao Chen, Lisha Li, and Joan Bruna. Supervised community detection with line graph neural networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- [8] Zhengdao Chen, Soledad Villar, Lei Chen, and Joan Bruna. On the equivalence between graph isomorphism testing and function approximation with GNNs. In *Advances in Neural Information Processing System (NeurIPS)*, pages 1–19, 2019.
- [9] Feng Dai and Yuan Xu. *Approximation Theory and Harmonic Analysis on Spheres and Balls*. Springer, 2013.
- [10] George Dasoulas, Ludovic Dos Santos, Kevin Scaman, and Aladin Virmaux. Coloring graph neural networks for node disambiguation. *IJCAI International Joint Conference on Artificial Intelligence*, 2021-January:2126–2132, 2020.
- [11] Gwendoline de Bie, Gabriel Peyré, and Marco Cuturi. Stochastic deep networks. In *International Conference on Machine Learning*, 2018.
- [12] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Advances in Neural Information and Processing Systems (NIPS)*, 2016.
- [13] Nima Dehmamy, Albert-László Barabási, and Rose Yu. Understanding the representation power of graph neural networks in learning graph topology. In *Advances in Neural Information and Processing Systems (NeurIPS)*, 2019.
- [14] Costas Efthimiou and Christopher Frye. *Spherical harmonics in p dimensions*. World Scientific, 2014.

- [15] Vikas K. Garg, Stefanie Jegelka, and Tommi Jaakkola. Generalization and representational limits of graph neural networks. In *International Conference on Machine Learning (ICML)*, pages 1–22, 2020.
- [16] Floris Geerts. The expressive power of kth-order invariant graph networks. *arXiv:2007.12035*, 2020.
- [17] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural Message Passing for Quantum Chemistry. In *International Conference on Machine Learning (ICML)*, pages 1–14, 2017.
- [18] William L. Hamilton. *Graph Representation Learning*. Morgan & Claypool, 2020.
- [19] Paul W Holland. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983.
- [20] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer Feedforward Networks are Universal Approximators. *Neural Networks*, 2:359–366, 1989.
- [21] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open Graph Benchmark: Datasets for Machine Learning on Graphs. *Neural Information Processing Systems (NeurIPS)*, 2020.
- [22] Nicolas Keriven, Alberto Bietti, and Samuel Vaiter. Convergence and Stability of Graph Convolutional Networks on Large Random Graphs. In *Advances in Neural Information and Processing Systems (NeurIPS)*, pages 1–26, 2020.
- [23] Nicolas Keriven and Gabriel Peyré. Universal Invariant and Equivariant Graph Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1–19, 2019.
- [24] Thomas N Kipf and Max Welling. Semi-Supervised Learning with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- [25] Jing Lei and Alessandro Rinaldo. Consistency of spectral clustering in stochastic block models. *Annals of Statistics*, 43(1):215–237, 2015.
- [26] Ron Levie, Michael M. Bronstein, and Gitta Kutyniok. Transferability of spectral graph convolutional neural networks. *arXiv:1907.12972*, pages 1–41, 2019.
- [27] Andreas Loukas. How hard is to distinguish graphs with graph neural networks? In *Advances in Neural Information and Processing Systems (NeurIPS)*, 2020.
- [28] Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *ICLR*, 2020.
- [29] László Lovász. Large networks and graph limits. *Colloquium Publications*, 60, 2012.
- [30] Abram Magnier, Mayank Baranwal, and Alfred O. Hero. The Power of Graph Convolutional Networks to Distinguish Random Graph Models. *arXiv:1910.12954*, pages 1–27, 2019.
- [31] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably Powerful Graph Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1–12, 2019.
- [32] Haggai Maron, Ethan Fetaya, Nimrod Segol, and Yaron Lipman. On the Universality of Invariant Networks. In *International Conference on Machine Learning (ICML)*, 2019.
- [33] Christine Leigh Myers Nickel. *Random Dot Product Graphs a Model for Social Networks*. PhD thesis, 2006.
- [34] Mathew Penrose. *Random Geometric Graphs*. Oxford University Press, 2008.
- [35] Allan Pinkus. Approximation theory of the MLP model in neural networks. *Acta Numerica*, 8(May):143–195, 1999.
- [36] Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.

- [37] Lorenzo Rosasco, Mikhail Belkin, and Ernesto De Vito. On learning with integral operators. *Journal of Machine Learning Research*, 11:905–934, 2010.
- [38] Luana Ruiz, Luiz F.O. Chamon, and Alejandro Ribeiro. Graphon neural networks and the transferability of graph neural networks. In *Advances in Neural Information and Processing Systems (NeurIPS)*, 2020.
- [39] Ryoma Sato, Makoto Yamada, and Hisashi Kashima. Random Features Strengthen Graph Neural Networks. *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*, 1:333–341, 2021.
- [40] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*. Cambridge University Press, 2018.
- [41] Clément Vignac, Andreas Loukas, and Pascal Frossard. Building powerful and equivariant graph neural networks with message-passing. In *Advances in Neural Information and Processing Systems (NeurIPS)*, 2020.
- [42] Ulrike Von Luxburg, Mikhail Belkin, and Olivier Bousquet. Consistency of spectral clustering. *Annals of Statistics*, 36(2):555–586, 2008.
- [43] B Yu Weisfeiler and A A Leman. The Reduction of a Graph to Canonical Form and the Algebra Which Appears Therein. *Nti*, 2(9):12–16, 1968.
- [44] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–21, 2020.
- [45] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How Powerful are Graph Neural Networks? In *ICLR*, pages 1–15, 2019.
- [46] Gilad Yehudai, Ethan Fetaya, Eli Meir, Gal Chechik, and Haggai Maron. On Size Generalization in Graph Neural Networks. *arXiv:2010.08853*, 2020.
- [47] Stephen J. Young and Edward R. Scheinerman. Random dot product graph models for social networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 4863 LNCS:138–149, 2007.
- [48] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J. Smola. Deep sets. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017-December(ii):3392–3402, 2017.
- [49] Aaron Zweig and Joan Bruna. A functional perspective on learning symmetric functions with neural networks. *arXiv*, pages 1–19, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) **See Sec. 6**
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#) **This is a theoretical work.**
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#) **In the supplementary material.**
3. If you ran experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes] As supplementary material.**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes] In App. E and the provided code.**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes] In Fig. 2.**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[N/A] Our experiments are purely illustrative and run on any machine.**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? **[N/A]**
 - (b) Did you mention the license of the assets? **[N/A]**
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

A Convergence

Let us start with some notations. Given a GNN Φ , we define some bounds on its parameters that will be used in the multiplicative constants of the theorem. Recall that the filters are written $h_{ij}^{(\ell)}(\lambda) = \sum_{k=0}^{\infty} \beta_{ijk}^{(\ell)} \lambda^k$. We define $B_k^{(\ell)} = \left(\beta_{ijk}^{(\ell)} \right)_{ji} \in \mathbb{R}^{d_{\ell+1} \times d_{\ell}}$ the matrix containing the order- k coefficients, and by $B_{k,|\cdot|}^{(\ell)} = \left(|\beta_{ijk}^{(\ell)}| \right)_{ji}$ the same matrix with absolute value on all coefficients. Recall that $\|\cdot\|$ is the operator norm for matrices. We define the following bounds:

$$\begin{aligned} H_2^{(\ell)} &= \sum_k \|B_k^{(\ell)}\| & H_{\partial,2}^{(\ell)} &= \sum_k \|B_k^{(\ell)}\| k \\ H_{\infty}^{(\ell)} &= \|B_{0,|\cdot|}^{(\ell)}\| + \sum_{k \geq 1} \|B_k^{(\ell)}\| & H_{\partial,\infty}^{(\ell)} &= \sum_k \|B_k^{(\ell)}\| k \sqrt{\log k} \end{aligned}$$

which all converge by our assumptions on the β_k . We may also denote H_2 by $H_{L^2(P)}$ for convenience but this quantity does not depend on P . Note that, only for H_{∞} , we use the spectral norm of the matrix $B_{0,|\cdot|}$ with non-negative coefficients, which is suboptimal compared to using B_0 . This is due to a part of our analysis where we do not operate in a Hilbert space but only in a Banach space $\mathcal{B}(\mathcal{X})$, see Lemma 4. We also define $\|b^{(\ell)}\| = \sqrt{\sum_j (b_j^{(\ell)})^2}$ to measure the norm of the bias.

Given $X = \{x_1, \dots, x_n\}$ and any dimension d , we denote by S_X the sampling operator acting on functions $f : \mathcal{X} \rightarrow \mathbb{R}^d$ defined by $S_X f \stackrel{\text{def.}}{=} [f(x_1), \dots, f(x_n)] \in \mathbb{R}^{n \times d}$. We have $\left\| \frac{1}{\sqrt{n}} S_X f \right\|_F \leq \|f\|_{\infty}$. Finally, given X and W , we define $W(X) \stackrel{\text{def.}}{=} (W(x_i, x_j))_{ij} \in \mathbb{R}^{n \times n}$, and remark that $\frac{W(X)}{n} \circ S_X = S_X \circ T_{W,X}$. In the deterministic edges case, the adjacency matrix A is directly $W(X)$. In the random edges case, A has expectation $W(X)$ (conditionally on X).

A.1 Convergence of GNNs: proof of Theorem 1

This proof is a variant of [22]. We prove Theorem 1 with the following error terms:

$$R_1(n) = \frac{C_1 \sqrt{d_{\mathcal{X}}} + C_2 \sqrt{\log \left(\frac{\sum_{\ell} d_{\ell}}{\rho} \right)}}{\sqrt{n}} \quad R_2(n) = \frac{C_{\nu} C_3}{\sqrt{\alpha_n n}} \quad R_3(n) = C_4 \sqrt{\frac{\log(1/\rho)}{n}}, \quad (11)$$

where R_2 is present in the random edges case and R_3 in the permutation-invariant case, and the following constants:

$$\begin{aligned} C &= L_g \prod_{\ell=0}^{M-1} H_2^{(\ell)} & C_1 &= \sum_{\ell=0}^{M-1} C^{(\ell)} H_{\partial,\infty}^{(\ell)} \\ C_2 &= C_1 (1 + D_{\mathcal{X}} L_W) & C_3 &= \sum_{\ell=0}^{M-1} C^{(\ell)} H_{\partial,2}^{(\ell)} & C_4 &= C^{(M)} \end{aligned}$$

where the MLP g is L_g -Lipschitz and $D_{\mathcal{X}} = \sup_{x, x' \in \mathcal{X}} m_{\mathcal{X}}(x, x')$ is the diameter of \mathcal{X} , and

$$C^{(\ell)} = L_g \left(\prod_{s=\ell+1}^{M-1} H_2^{(s)} \right) \left(C_f \prod_{s=0}^{\ell-1} H_{\infty}^{(s)} + \sum_{s=0}^{\ell-1} \|b^{(s)}\| \prod_{p=s+1}^{\ell-1} H_{\infty}^{(p)} \right)$$

with the conventions that an empty product is 1 and an empty sum is 0.

We begin the proof by the equivariant case, the invariant case will simply use an additional concentration inequality. We have

$$\begin{aligned} \text{MSE}_X \left(\Phi_A(Z^{(0)}), \Phi_{W,P}(f) \right) &= \frac{1}{\sqrt{n}} \left\| \Phi_A(Z^{(0)}) - S_X \Phi_{W,P}(f^{(0)}) \right\|_F \\ &\leq L_g \frac{1}{\sqrt{n}} \left\| Z^{(M)} - S_X f^{(M)} \right\|_F. \end{aligned}$$

We therefore seek to bound that last term. Define the notation

$$\Delta^{(\ell)} = \frac{1}{\sqrt{n}} \sqrt{\sum_j \left\| \sum_i \left(h_{ij}^{(\ell)} \left(\frac{A}{n} \right) S_X f_i^{(\ell)} - S_X h_{ij}^{(\ell)}(T_{W,P}) f_i^{(\ell)} \right) \right\|^2}.$$

Then, using the Lipschitzness of ρ , Lemma 4 with $\|A/n\| \leq 1$, and the fact that $S_X \circ \rho = \rho \circ S_X$, we have

$$\begin{aligned} & \left\| Z^{(\ell+1)} - S_X f^{(\ell+1)} \right\|_F \\ &= \left(\sum_j \left\| \rho \left(\sum_{i=1}^{d_\ell} h_{ij}^{(\ell)} \left(\frac{A}{n} \right) Z_{:,i}^{(\ell)} + b_j^{(\ell)} \mathbf{1}_n \right) - S_X \rho \left(\sum_{i=1}^{d_\ell} h_{ij}^{(\ell)}(T_{W,P}) f_i^{(\ell)} + b_j^{(\ell)} \right) \right\|^2 \right)^{\frac{1}{2}} \\ &= \left(\sum_j \left\| \rho \left(\sum_{i=1}^{d_\ell} h_{ij}^{(\ell)} \left(\frac{A}{n} \right) z_i^{(\ell)} + b_j^{(\ell)} \mathbf{1}_n \right) - \rho \left(S_X \left(\sum_{i=1}^{d_\ell} h_{ij}^{(\ell)}(T_{W,P}) f_i^{(\ell)} + b_j^{(\ell)} \right) \right) \right\|^2 \right)^{\frac{1}{2}} \\ &\leq \left(\sum_j \left\| \sum_{i=1}^{d_\ell} h_{ij}^{(\ell)} \left(\frac{A}{n} \right) Z_{:,i}^{(\ell)} - S_X h_{ij}^{(\ell)}(T_{W,P}) f_i^{(\ell)} \right\|^2 \right)^{\frac{1}{2}} \\ &\leq \left(\sum_j \left\| \sum_{i=1}^{d_\ell} h_{ij}^{(\ell)} \left(\frac{A}{n} \right) (Z_{:,i}^{(\ell)} - S_X f_i^{(\ell)}) \right\|^2 \right)^{\frac{1}{2}} \\ &\quad + \left(\sum_j \left\| \sum_{i=1}^{d_\ell} h_{ij}^{(\ell)} \left(\frac{A}{n} \right) S_X f_i^{(\ell)} - S_X h_{ij}^{(\ell)}(T_{W,P}) f_i^{(\ell)} \right\|^2 \right)^{\frac{1}{2}} \\ &\leq H_2^{(\ell)} \left\| Z^{(\ell)} - S_X f^{(\ell)} \right\|_F + \sqrt{n} \Delta^{(\ell)}. \end{aligned}$$

A recursion shows that, for all $Z^{(0)}$:

$$\begin{aligned} \text{MSE}_X \left(\Phi_A(Z^{(0)}), \Phi_{W,P}(f) \right) &\leq L_g \sum_{\ell=0}^{M-1} \Delta^{(\ell)} \prod_{s=\ell+1}^{M-1} H_2^{(s)} \\ &\quad + L_g \left(\prod_{\ell=0}^{M-1} H_2^{(\ell)} \right) \text{MSE}_X(Z^{(0)}, f^{(0)}). \end{aligned} \quad (12)$$

We now bound all $\Delta^{(\ell)}$ with high probability. Recall that $\frac{W(X)}{n} \circ S_X = S_X \circ T_{W,X}$, and that we have $\left\| \frac{S_X}{\sqrt{n}} f \right\| \leq \|f\|_\infty$ and $\|T_{W,X}\|_\infty \leq 1$. By Lemma 4 we have

$$\begin{aligned} \Delta^{(\ell)} &\leq \sqrt{\sum_j \left\| \sum_i \left(h_{ij}^{(\ell)} \left(\frac{A}{n} \right) - h_{ij}^{(\ell)} \left(\frac{W(X)}{n} \right) \right) \frac{S_X}{\sqrt{n}} f_i^{(\ell)} \right\|^2} \\ &\quad + \sqrt{\sum_j \left\| \sum_i \frac{S_X}{\sqrt{n}} \left(h_{ij}^{(\ell)}(T_{W,X}) - h_{ij}^{(\ell)}(T_{W,P}) \right) f_i^{(\ell)} \right\|^2} \\ &\leq H_{\partial,2}^{(\ell)} \left\| \frac{A - W(X)}{n} \right\| \left\| f^{(\ell)} \right\|_\infty \\ &\quad + \sum_k \|B_k\| \sqrt{\sum_i \left(\sum_{p=0}^{k-1} \left\| (T_{W,X} - T_{W,P}) T_{W,P}^{k-1-p} f_i^{(\ell)} \right\|_\infty \right)^2}. \end{aligned} \quad (13)$$

The first term in (13) is 0 in the deterministic edges case. In the random edges case, it is handled with a recent concentration inequality for Bernoulli matrices [25], recalled in Theorem 5 in App. D. Since $\alpha_n \gtrsim \frac{\log n}{n}$, for any ν , there is a constant C_ν such that, with probability $1 - n^{-\nu}$ on the random edges (conditionally on X), $\left\| \frac{A-W(X)}{n} \right\| \leq \frac{C_\nu}{\sqrt{\alpha_n n}}$. By the law of total probability, it is valid with joint probability $1 - n^{-\nu}$ on both X and the random edges.

We now bound the second term in (13). Define $\rho_k = \frac{C\rho}{(k+1)^2 \sum_\ell d_\ell}$ with C such that $\sum_{k\ell} d_\ell \rho_k = \rho/4$ (even when the filters are not of finite order). Using an application of Dudley's inequality detailed in Lemma 3, applied with $U(x, y) = W(x, y)f(y)$ which is bounded by $\|f\|_\infty$ and has Lipschitz constant $L_W \|f\|_\infty$ in the first variable, and a union bound, we obtain with probability $1 - \rho/4$ that: for all i, ℓ, k , we have

$$\left\| (T_{W,X} - T_{W,P}) T_{W,P}^k f_i^{(\ell)} \right\|_\infty \lesssim \frac{1}{\sqrt{n}} \|f_i^{(\ell)}\|_\infty \left(\sqrt{d_X} + (1 + D_X L_W) \sqrt{\log \rho_k^{-1}} \right).$$

Coming back to the second term of (13), with probability $1 - \rho/4$:

$$\begin{aligned} & \sum_k \|B_k\| \sqrt{\sum_i \left(\sum_{p=0}^{k-1} \left(\left\| (T_{W,X} - T_{W,P}) T_{W,P}^{k-1-p} f_i^{(\ell)} \right\|_\infty \right)^2 \right)} \\ & \lesssim \frac{\left(\sqrt{d_X} + (1 + D_X L_W) \sqrt{\log \frac{\sum_\ell d_\ell}{\rho}} \right)}{\sqrt{n}} \sum_k \|B_k\| k \sqrt{\log k} \sqrt{\sum_i \left\| f_i^{(\ell)} \right\|_\infty^2} \\ & \leq \frac{\left(\sqrt{d_X} + (1 + D_X L_W) \sqrt{\log \frac{\sum_\ell d_\ell}{\rho}} \right)}{\sqrt{n}} H_{\partial, \infty}^{(\ell)} \|f^{(\ell)}\|_\infty. \end{aligned}$$

At the end of the day we obtain that with probability $1 - \rho$, for all ℓ :

$$\Delta^{(\ell)} \propto \|f^{(\ell)}\|_\infty \left(\frac{H_{\partial, 2}^{(\ell)}}{\sqrt{\alpha_n n}} + \frac{H_{\partial, \infty}^{(\ell)} \left(\sqrt{d} + (1 + D_X L_W) \sqrt{\log \frac{\sum_\ell d_\ell}{\rho}} \right)}{\sqrt{n}} \right).$$

We then use Lemma 5 to bound $\|f^{(\ell)}\|_\infty$ and conclude.

For the invariant case, we have

$$\begin{aligned} \left\| \bar{\Phi}_A(Z^{(0)}) - \bar{\Phi}_{W,P}(f^{(0)}) \right\| & \leq \text{MSE}_X(\Phi_A(Z^{(0)}), \Phi_{W,P}(f^{(0)})) \\ & \quad + L_g \left\| \frac{1}{n} \sum_{i=1}^n f^{(M)}(x_i) - \int f^{(M)}(x) dP(x) \right\| \end{aligned}$$

We use a vector Hoeffding's inequality [36, Lemma 4] and a bound on $\|f^{(M)}\|_\infty$ (Lemma 5) to conclude.

A.2 Convergence of SGNNs

We prove Theorem 2 with the same form of error terms (11) where R_i is replaced by R'_i with modified multiplicative constants C'_i . Here we will have:

$$\begin{aligned}
C' &= DL_g \prod_{\ell=0}^{M-1} H_2^{(\ell)} \\
C'_1 &= DL_\Phi + \sum_{\ell=0}^{M-1} H'_{\partial,\infty}{}^{(\ell)} C'^{(\ell)} + \sum_{\ell=0}^{M-1} H_{\partial,\infty}^{(\ell)} C^{(\ell)} \\
C'_2 &= (1 + D_{\mathcal{X}} L_W) \sum_{\ell=0}^{M-1} H'_{\partial,\infty}{}^{(\ell)} C'^{(\ell)} + DL_\Phi D_{\mathcal{X}} + DC_\Phi + \sum_{\ell=0}^{M-1} H_{\partial,\infty}^{(\ell)} (C^{(\ell)} + D_{\mathcal{X}} L^{(\ell)}) \\
C'_3 &= \sum_{\ell=0}^{M-1} H'_{\partial,2}{}^{(\ell)} C'^{(\ell)} + H_{\partial,2}^{(\ell)} C^{(\ell)}, \quad C'_4 = C'^{(M)}
\end{aligned}$$

where $H_\star'^{(\ell)}$ is like $H_\star^{(\ell)}$ but for the weights in Φ' , the final-layer MLP of Φ' is denoted by g' with a Lipschitz constant $L_{g'}$, and:

$$\begin{aligned}
D &= L_{g'} \prod_{\ell=0}^{M-1} H_2'^{(\ell)} \\
C'^{(\ell)} &= L_{g'} \left(\prod_{s=\ell+1}^{M-1} H_2'^{(s)} \right) \left(C_\Phi \prod_{s=0}^{\ell-1} H_\infty'^{(s)} + \sum_{s=0}^{\ell-1} \|b'^{(s)}\| \prod_{p=s+1}^{\ell-1} H_\infty'^{(p)} \right) \\
C^{(\ell)} &= DL_g \left(\prod_{s=\ell+1}^{M-1} H_2^{(s)} \right) \tilde{C}_\infty^{(\ell)} \\
L^{(\ell)} &= DL_g \left(\prod_{s=\ell+1}^{M-1} H_2^{(s)} \right) \left(L_W \tilde{C}_\infty^{(\ell)} + \sqrt{d_\ell} L_\eta \prod_{s=0}^{\ell-1} H_\infty^{(s)} \right) \\
C_\Phi &= \|g(0)\| + L_g \tilde{C}_\infty^{(M)} \\
L_\Phi &= L_g \left(L_\eta \prod_{\ell=0}^{M-1} \|B_0^{(\ell)}\| \right) + L_W \sum_{\ell=0}^{M-1} \left(\prod_{s=\ell+1}^{M-1} \|B_0^{(s)}\| \right) \tilde{C}_2^{(\ell)}
\end{aligned}$$

with

$$\tilde{C}_\star^{(\ell)} = C_\eta \prod_{s=0}^{\ell-1} H_\star^{(s)} + \sum_{s=0}^{\ell-1} \|b^{(s)}\| \prod_{p=s+1}^{\ell-1} H_\star^{(p)} \quad \text{for } \star \in \{2, \infty\}.$$

We start by applying Theorem 1 on the outer GNN Φ' . Since the result is uniformly valid over all input of the GNN $Z^{(0)}$ with probability $1 - \rho$:

$$\text{MSE}_X(\Psi_A, \Psi_{W,P}) \leq DMSE_X \left(\frac{1}{n} \sum_q \Phi_A(E_q(A)), \int \Phi_{W,P}(\eta(\cdot, x)) dP(x) \right) + R'(n) \quad (14)$$

where, from Theorem 1, $D = L_{g'} \prod_{\ell=0}^{M-1} H_2'^{(\ell)}$ is C_1 but for the weights in Φ' , and $R'(n)$ is the error term formed by summing various $R'_i(n)$, taking into account that by Lemma 5 the function inputed in Φ' is bounded by C_Φ .

We must therefore bound the first term in (14). We write

$$\begin{aligned}
& \text{MSE}_X \left(\frac{1}{n} \sum_q \Phi_A(E_q(A)), \int \Phi_{W,P}(\eta(\cdot, x)) dP(x) \right) \\
& \leq \text{MSE}_X \left(\frac{1}{n} \sum_q \Phi_A(E_q(A)), \frac{1}{n} \sum_q \Phi_{W,P}(\eta(\cdot, x_q)) \right) \\
& \quad + \text{MSE}_X \left(\frac{1}{n} \sum_q \Phi_{W,P}(\eta(\cdot, x_q)), \int \Phi_{W,P}(\eta(\cdot, x)) dP(x) \right) \\
& \leq \sup_q \text{MSE}_X (\Phi_A(E_q(A)), \Phi_{W,P}(\eta(\cdot, x_q))) \\
& \quad + \left\| \frac{1}{n} \sum_q \Phi_{W,P}(\eta(\cdot, x_q)) - \int \Phi_{W,P}(\eta(\cdot, x)) dP(x) \right\|_\infty. \tag{15}
\end{aligned}$$

Let us start with the second term. By Lemma 5 and the Lipschitzness of g , $U(x, y) \stackrel{\text{def.}}{=} \Phi_{W,P}(\eta(\cdot, y))(x)$ is C_Φ -bounded and L_Φ -Lipschitz with respect to x .

Hence, applying Lemma 3: with probability $1 - \rho$,

$$\left\| \frac{1}{n} \sum_q \Phi_{W,P}(\eta(\cdot, x_q)) - \int \Phi_{W,P}(\eta(\cdot, x)) dP(x) \right\|_\infty \lesssim \frac{L_\Phi \sqrt{d_{\mathcal{X}}} + (L_\Phi D_{\mathcal{X}} + C_\Phi) \sqrt{\log(1/\rho)}}{\sqrt{n}}. \tag{16}$$

For the first term in (15), we introduce some notations. We denote by $f_i^{(\ell)} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ the bivariate function propagated at each layer of the inner part of the c-SGNN, as:

$$f_0^{(0)} = \eta \quad f_j^{(\ell+1)} = \rho \left(\sum_i h_{ij}^{(\ell)} (T_{W,P}) f_i^{(\ell)} + b_j^{(\ell)} \right) \tag{17}$$

where $T_{W,P}$ is here to be understood as an operator on $\mathcal{C}(\mathcal{X} \times \mathcal{X})$ defined by $T_{W,P}[f](x, y) = \int W(x, z) f(z, y) dP(z)$. With these notations, $\Phi_{W,P}(\eta(\cdot, x_q)) = g(f^{(M)}(\cdot, x_q))$. Note that we still have $\|T_{W,P}\|_\infty \leq 1$ for this version. We perform the computation as in the proof of Theorem 1 in (12) to obtain:

$$\begin{aligned}
& \sup_q \text{MSE}_X (\Phi_A(E_q(A)), \Phi_{W,P}(\eta(\cdot, x_q))) \\
& \leq L_g \sum_{\ell=0}^{M-1} \sup_q \Delta_q^{(\ell)} \prod_{s=\ell+1}^{M-1} H_2^{(s)} + L_g \left(\prod_{\ell=0}^{M-1} H_2^{(\ell)} \right) \sup_q \text{MSE}_X (E_q(A), \eta(\cdot, x_q)) \tag{18}
\end{aligned}$$

with

$$\Delta_q^{(\ell)} = \frac{1}{\sqrt{n}} \sqrt{\sum_j \left\| \sum_i \left(h_{ij}^{(\ell)} \left(\frac{A}{n} \right) S_X f_i^{(\ell)}(\cdot, x_q) - S_X h_{ij}^{(\ell)} (T_{W,P}) [f_i^{(\ell)}(\cdot, x_q)] \right) \right\|_\infty^2}. \tag{19}$$

Then, again we decompose

$$\begin{aligned}
\sup_q \Delta_q^{(\ell)} & \leq H_{\partial,2}^{(\ell)} \left\| \frac{A - W(X)}{n} \right\|_\infty \|f^{(\ell)}\|_\infty \\
& \quad + \sum_k \|B_k\| \sqrt{\sum_i \left(\sum_{p=0}^{k-1} \left\| (T_{W,X} - T_{W,P}) T_{W,P}^{k-1-p} [f_i^{(\ell)}] \right\|_\infty \right)^2} \tag{20}
\end{aligned}$$

where we recall here that $f_i^{(\ell)}$ is a bivariate function.

Again, the first term is 0 in the deterministic edges case, and otherwise by Theorem 5 we have $\left\| \frac{A-W(X)}{n} \right\| \leq C_\nu / \sqrt{\alpha_n n}$ with probability $1 - n^{-\nu}$, and by Lemma 6 we have

$$\|f^{(\ell)}\|_\infty \leq C_\eta \prod_{s=0}^{\ell-1} H_\infty^{(s)} + \sum_{s=0}^{\ell-1} \|b^{(s)}\| \prod_{p=s+1}^{\ell-1} H_\infty^{(p)}.$$

Fix k, ℓ, i for now. We will apply Lemma 3 with $U : (\mathcal{X} \times \mathcal{X}) \times \mathcal{X} \rightarrow \mathbb{R}$ defined as $U((x, x'), y) = W(x, y)f(y, x')$ for $f(y, x') = T_{W,P}^k[f_i^{(\ell)}(\cdot, x')](y)$. Since $\|T_{W,P}\|_\infty \leq 1$ we have $\|f\|_\infty \leq \|f_i^{(\ell)}\|_\infty$. Then,

$$\begin{aligned} \left\| T_{W,P}^k[f_i^{(\ell)}(\cdot, x')] - T_{W,P}^k[f_i^{(\ell)}(\cdot, x'')] \right\|_\infty &= \left\| T_{W,P}^k[f_i^{(\ell)}(\cdot, x') - f_i^{(\ell)}(\cdot, x'')] \right\|_\infty \\ &\leq \left\| f_i^{(\ell)}(\cdot, x') - f_i^{(\ell)}(\cdot, x'') \right\|_\infty \\ &\leq L_\eta m_{\mathcal{X}}(x, x') \prod_{s=0}^{\ell-1} H_\infty^{(s)} \end{aligned}$$

by Lemma 6. Hence U is bounded by $\|f_i^{(\ell)}\|_\infty$ and $L_i^{(\ell)}$ -Lipschitz with respect to (x, x') , with

$$L_i^{(\ell)} = L_W \left\| f_i^{(\ell)} \right\|_\infty + L_\eta \prod_{s=0}^{\ell-1} H_\infty^{(s)}. \quad (21)$$

Finally, note that $\mathcal{X} \times \mathcal{X}$ is compact with covering numbers proportional to ε^{-2d} . Hence by Lemma 3 and a union bound, again defining ρ_k as in the proof of Theorem 1 such that $\sum_{ik\ell} \rho_k = \rho$: with probability $1 - \rho$, we have simultaneously for all i, k, ℓ :

$$\begin{aligned} \sup_x \left\| (T_{W,X} - T_{W,P}) T_{W,P}^k f_i^{(\ell)}(\cdot, x) \right\|_\infty \\ \lesssim \frac{1}{\sqrt{n}} \left(\left\| f_i^{(\ell)} \right\|_\infty \sqrt{d_{\mathcal{X}}} + \left(\left\| f_i^{(\ell)} \right\|_\infty + D_{\mathcal{X}} L_i^{(\ell)} \right) \sqrt{\log \rho_k^{-1}} \right). \end{aligned}$$

Hence, as in the previous proof:

$$\begin{aligned} \sum_k \|B_k\| \sqrt{\sum_i \left(\sum_{p=0}^{k-1} \left(\left\| (T_{W,X} - T_{W,P}) T_{W,P}^{k-1-p} f_i^{(\ell)} \right\|_\infty \right)^2 \right)} \\ \leq \frac{\left(\left\| f^{(\ell)} \right\|_\infty \sqrt{d_{\mathcal{X}}} + \left(\left\| f^{(\ell)} \right\|_\infty + D_{\mathcal{X}} L^{(\ell)} \right) \sqrt{\log \frac{\sum_\ell d_\ell}{\rho}} \right) H_{\partial, \infty}^{(\ell)}}{\sqrt{n}} \end{aligned}$$

where $L^{(\ell)} = (\sum_i (L_i^{(\ell)})^2)^{\frac{1}{2}}$.

A.3 Proof of Prop. 1

The error can be written as

$$\begin{aligned} \frac{1}{\sqrt{n}} \left\| A^2 e_q / n - S_X T_{W,P}(W(\cdot, x_q)) \right\|_2 &\leq \frac{1}{\sqrt{n}} \left\| A^2 e_q / n - W(X)^2 e_q / n \right\|_2 \\ &\quad + \left\| (T_{W,X} - T_{W,P}) W(\cdot, x_q) \right\|_\infty. \end{aligned}$$

Using chaining as in the previous section, we have

$$\sup_x \left\| (T_{W,X} - T_{W,P}) W(\cdot, x) \right\|_\infty \lesssim \frac{1}{\sqrt{n}} \left(\sqrt{d_{\mathcal{X}}} + (1 + D_{\mathcal{X}} L_W) \sqrt{\log 1/\rho} \right).$$

The first term is 0 in the deterministic edges case, and otherwise:

$$\begin{aligned} \frac{1}{\sqrt{n}} \|A^2 e_q/n - W(X)^2 e_q/n\|_2 &= \left(\frac{1}{n} \sum_i \left(\frac{1}{n} \sum_j a_{ij} a_{jq} - \frac{1}{n} \sum_j W(x_i, x_j) W(x_j, x_q) \right)^2 \right)^{\frac{1}{2}} \\ &\leq \left(\frac{1}{n} \sum_{i \neq q} \left(\frac{1}{n} \sum_j a_{ij} a_{jq} - W(x_i, x_j) W(x_j, x_q) \right)^2 \right)^{\frac{1}{2}} \\ &\quad + \frac{1}{\sqrt{n}} \left(\frac{1}{n} \sum_j a_{jq}^2 - \frac{1}{n} \sum_j W(x_j, x_q)^2 \right)^2. \end{aligned}$$

Now, by Bernstein inequality with

$$\text{Var}(a_{ij} a_{jq}) \leq \mathbb{E}(a_{ij}^2 a_{jq}^2) = \alpha_n^{-2} \mathbb{E}(a_{ij} a_{jq}) = \alpha_n^{-2} W(x_i, x_j) W(x_j, x_q)$$

for $i \neq q$ and a union bound, with proba $1 - \delta$, we have:

$$\left| \frac{1}{n} \sum_j a_{ij} a_{jq} - \frac{1}{n} \sum_j W(x_i, x_j) W(x_j, x_q) \right| \lesssim \frac{\sqrt{\log(n/\delta)}}{\alpha_n \sqrt{n}} \text{ for all } q \text{ and } i \neq q.$$

Since $\left| \frac{1}{n} \sum_j a_{jq}^2 - \frac{1}{n} \sum_j W(x_j, x_q)^2 \right| \leq 2$, we have

$$\frac{1}{\sqrt{n}} \sup_q \|A^2 e_q/n - S_X T_W(W(\cdot, x_q))\|_2 \lesssim \frac{\sqrt{\log(n/\delta)}}{\alpha_n \sqrt{n}} + \frac{\sqrt{d_X} + (1 + D_X L_W) \sqrt{\log 1/\rho}}{\sqrt{n}} \quad (22)$$

B Approximation power: invariant case

B.1 Application of Stone-Weierstrass

Proof of Prop. 2. We do the proof for cSGNNs, it is exactly similar for cGNNs.

This is a direct application of Lemma 1: for any two cSGNNs $\bar{\Psi} : \mathcal{W} \times \mathcal{P} \rightarrow \mathbb{R}^d$, $\bar{\Psi}' : \mathcal{W} \times \mathcal{P} \rightarrow \mathbb{R}^{d'}$, their concatenation $[\bar{\Psi}, \bar{\Psi}'] : \mathcal{W} \times \mathcal{P} \rightarrow \mathbb{R}^{d+d'}$ is also a cSGNN (if they do not use the same input transforms η, η' , one can concatenate $\eta'' = [\eta, \eta']$), and for any MLP g , $g \circ \bar{\Psi}$ is also a cSGNN.

One must just check that cSGNNs are continuous with respect to $\|\cdot\|_\infty + \|\cdot\|_{\text{TV}}$ on $\mathcal{W} \times \mathcal{P}$:

- $(W, P) \mapsto \eta$ is continuous by assumption ;
- for any $f_{W,P} \in \mathcal{C}(\mathcal{X}, \mathbb{R}^d)$ continuously indexed by (W, P) ,

$$\begin{aligned} \|T_{W,P}[f_{W,P}] - T_{W',P'}[f_{W',P'}]\|_\infty &\leq \|f_{W,P}\|_\infty (\|W - W'\|_\infty + \|P - P'\|_{\text{TV}}) \\ &\quad + \|f_{W,P} - f_{W',P'}\|_\infty \end{aligned}$$

and similarly for $\left\| \int f_{W,P} dP - \int f_{W',P'} dP' \right\|$;

- the non-linearity ρ is Lipschitz.

□

B.2 cSGNNs are more powerful than cGNNs

Proof of Theorem 3. By construction, cGNNs are included in cSGNNs, since one can take $\Phi = 0$ as the input GNN before pooling in (6).

To prove strict inclusion, we will construct two models $(W, P), (W', P')$ such that, for any cGNN we have $\bar{\Phi}_{W,P} = \bar{\Phi}_{W',P'}$, but there exists a cSGNN such that $\bar{\Psi}_{W,P} \neq \bar{\Psi}_{W',P'}$. We do the proof in

the random edges case with two-hop input filtering $\eta_{W,P} = T_{W,P}(W)$, since such cSGNNs can of course also be constructed in the deterministic edges case.

Since \mathcal{X} is not a singleton, one can single out two arbitrary elements x, x' and take P as a sum of two Diracs over them, which is equivalent to considering that $\mathcal{X} = \{x, x'\}$ (since any invariant architecture involves a final integration by P , it is useless to consider W outside of the support of P). This results in a two-community SBM, for which P can be represented as a 2-vector on the simplex and W as a 2-by-2 symmetric matrix. We then consider a family of SBMs indexed by $\gamma \in [0, 1]$:

$$P = \begin{pmatrix} 1/3 \\ 2/3 \end{pmatrix}, \quad W_\gamma = \begin{pmatrix} \gamma & \frac{1-\gamma}{2} \\ \frac{1-\gamma}{2} & \frac{1+\gamma}{4} \end{pmatrix}.$$

It is not hard to see that $T_{W_\gamma,P}[1] = 1/3 \cdot 1$ for any γ . Therefore, for any cGNN Φ , the function propagated inside its layers is always constant, and does not depend on γ . That is, $\bar{\Phi}_{W_\gamma,P} = \bar{\Phi}_{W_0,P}$ for any γ . On the other hand, consider the following SGNN:

$$\bar{\Psi}_{W,P} = \int_x \int_y f(T_W(W(\cdot, y))) dP(y) dP(x)$$

where f is an MLP. By the universality theorem, f can approximate $x \rightarrow x^2$, for which we obtain:

$$\bar{\Psi}_{W_\gamma,P} \approx 1/16 * \gamma^4 - 1/12 * \gamma^3 + 1/24 * \gamma^2 - 1/108 * \gamma + 17/1296.$$

This is not a constant function, so we can always find γ, γ' such that $\bar{\Psi}_{W_\gamma,P} \neq \bar{\Psi}_{W_{\gamma'},P}$, which concludes the proof. \square

B.3 SBMs

Proof of Prop. 3. We apply Prop. 2. We fix P as an incoherent vector in the k -simplex, and define $\mathcal{M} = \{(W, P) : W \in S_k([0, 1])\}$ which is indeed compact. It therefore suffices to show that cSGNNs separates points in \mathcal{M} .

We proceed by contraposition: assume that W, W' are such that $\bar{\Psi}_{W,P} = \bar{\Psi}_{W',P}$ for any cSGNN Ψ . We must show that necessarily $W = W'$. We look at cSGNNs of the form

$$\begin{aligned} \Psi_{W,P} &= \int f_1 \left(\int f_0(W(x, y)) dP(y) \right) dP(x) \\ &= \sum_i P_i f_1 \left(\sum_j P_j f_0(W_{ij}) \right) = \sum_i P_i f_1 \left(\sum_j P_j f_0(W'_{ij}) \right) \end{aligned}$$

where f_0, f_1 are MLPs. By the universality theorem, they can approximate any continuous function. Pick any f_0 . Then f_1 can be chosen as to take only values in $\{0, 1\}$ on the discrete set $\{\sum_j P_j f_0(W_{ij}), \sum_j P_j f_0(W'_{ij})\}_i$ of size $2K$. Moreover, if there was an index i_0 such that $\sum_j P_j f_0(W_{i_0j}) \neq \sum_j P_j f_0(W'_{i_0j})$, f_1 can be chosen to give different values on them. Then, defining $s_i = f_1(\sum_j P_j f_0(W_{ij})) - f_1(\sum_j P_j f_0(W'_{ij})) \in \{-1, 0, 1\}$, we have both $s_{i_0} \neq 0$ and

$$\sum_i P_i s_i = \sum_i P_i \left(f_1 \left(\sum_j P_j f_0(W_{ij}) \right) - f_1 \left(\sum_j P_j f_0(W'_{ij}) \right) \right) = 0$$

which contradicts the incoherence of P . So, for all f_0 and i , we have $\sum_j P_j f_0(W_{ij}) = \sum_j P_j f_0(W'_{ij})$. By the exact same reasoning on f_1 , we obtain that for all i, j , $W_{ij} = W'_{ij}$, which concludes the proof.

For the failure of c-GNNs, the proof is immediate using the example SBM in the proof of Theorem 3, since $P = [1/3, 2/3]$ is indeed incoherent. \square

B.4 Decomposed kernel

Proof of Prop. 4. Applying Prop. 2 with $\mathcal{M} = \{W\} \times \tilde{\mathcal{P}}$, it suffices to show that cSGNNs separate the distributions in $\tilde{\mathcal{P}}$. By contraposition, assume that $P, P' \in \tilde{\mathcal{P}}$ are such that $\bar{\Psi}_{W,P} = \bar{\Psi}_{W,P'}$ for any cSGNN, and we want to prove that necessarily $P = P'$.

We look at cSGNN of the form $P \mapsto \int f_1(\int f_0(W(x, y))dP(y))dP(x)$, where f_0, f_1 are MLPs, that can approximate any continuous functions by the universality theorem. Since u is continuous and injective, it is well-known that it has a continuous inverse on its image. Hence f_0 can be chosen to approximate $f_0 \approx u^{-1}$. By choosing f_1 to approximate $x \rightarrow x^k$, we obtain that:

$$\int (v(x) + \mathbb{E}_P v)^k dP(x) = \int (v(x) + \mathbb{E}_{P'} v)^k dP'(x).$$

Taking $k = 1$ we obtain that $\mathbb{E}_P v = \mathbb{E}_{P'} v$, and by an easy recursion we have $\mathbb{E}_P v^k = \mathbb{E}_{P'} v^k$ for all k . Since v is invertible and polynomial functions are universal approximators on compacts one can write $v^{-1}(x) = \sum_k a_k x^k$ and $x = \sum_k a_k v(x)^k$, such that $\mathbb{E}_P X^k = \mathbb{E}_{P'} X^k$. Again, by the universality of polynomial functions, $\mathbb{E}_P f = \mathbb{E}_{P'} f$ for any continuous function, which is well-known to be equivalent to $P = P'$ and concludes the proof. \square

B.5 Radial kernel

Proof of Prop. 5. We proceed as in the proof of Prop. 4 above: assuming $P, P' \in \tilde{\mathcal{P}}$ are such that $\bar{\Psi}_{W,P} = \bar{\Psi}_{W,P'}$ for any cSGNN, we want to prove that necessarily $P = P'$. We look at cSGNN of the form $P \mapsto \int f_1(\int f_0(W(x, y))dP(y))dP(x)$, where f_0, f_1 are MLPs. Since w is injective f_0 can approximate $(x \rightarrow x^2) \circ w^{-1}$. By choosing f_1 to approximate $x \rightarrow x^k$, and since P, P' are centered we obtain

$$\int (x^2 + \mathbb{E}_P X^2)^k dP(x) = \int (x^2 + \mathbb{E}_{P'} X^2)^k dP'(x).$$

Taking $k = 1$ we have $\mathbb{E}_P X^2 = \mathbb{E}_{P'} X^2$, and by an easy recursion $\mathbb{E}_P X^{2k} = \mathbb{E}_{P'} X^{2k}$ for all k . Since P, P' have 0 odd-order moments, $\mathbb{E}_P X^k = \mathbb{E}_{P'} X^k$ for all k , from which we can conclude $P = P'$ as in the previous proof. \square

C Approximation power: equivariant case

C.1 Application of Stone-Weierstrass

Proof of Prop. 6. As the proof of Prop. 2, this is a direct application of Lemma 1: the set of cSGNNs is closed by concatenation and composition with an MLP, \mathcal{X} is compact, and any equivariant cSGNN in continuous since by assumption $W, \eta_{W,P}$ and ρ are. \square

C.2 cSGNNs are more powerful than cGNNs

Proof of Theorem 4. As in the proof of Theorem 3 in App. B.2, non-strict inclusion is immediate. To prove strict inclusion, as in the proof of Theorem 3 we consider again the same 2-community SBM but for $\gamma = 1/2$:

$$P = \begin{pmatrix} 1/3 \\ 2/3 \end{pmatrix}, \quad W = \begin{pmatrix} 1/2 & 1/4 \\ 1/4 & 3/8 \end{pmatrix}.$$

Again any c-GNN would return a constant function $\Phi_{W,P}(1) = \Phi_{W,P}(2)$, while if we consider the following c-SGNN for two-hop filtering:

$$\Psi_{W,P} = \int f(T_W(W(\cdot, y)))dP(y)$$

with f an MLP that approximates $x \rightarrow x^2$, we obtain $\Psi_{W,P}(1) \approx 1/8$ and $\Psi_{W,P}(2) \approx 11/96$, hence a non-constant function. \square

C.3 SBMs

Proof of Prop. 7. We treat the two-hop filtering case, since they are included in one-hop architectures. By Prop. 6, we must prove the separation of elements of \mathcal{X} , which here are discrete community labels $\mathcal{X} = \{1, \dots, K\}$. Fix $P \in \Delta^{K-1}$ incoherent and $W \in S_K$ invertible. Assume that k, k' are two

communities such that $\Psi_{W,P}(k) = \Psi_{W,P}(k')$ for all Ψ with two-hop input filtering. We want to show that necessarily $k = k'$. By assumption, we have:

$$\sum_i f \left(\sum_j W_{kj} W_{ji} P_j \right) P_i = \sum_i f \left(\sum_j W_{k'j} W_{ji} P_j \right) P_i$$

for any MLP f . As in the proof of Prop. 3 in App. B.3, f can approximate a function that is $\{0, 1\}$ -valued on its inputs such that, if there is an index i such that $\sum_j W_{kj} W_{ji} P_j \neq \sum_j W_{k'j} W_{ji} P_j$, then the incoherency of P is contradicted. Hence, for all i , $\sum_j W_{kj} W_{ji} P_j = \sum_j W_{k'j} W_{ji} P_j$, or in other words:

$$W \cdot (P \odot (W_{k,:} - W_{k',:})) = 0.$$

Since W is invertible and P has only non-zero coordinates (by incoherency), we obtain $W_{k,:} = W_{k',:}$. Since W is invertible it has necessarily distinct columns, so $k = k'$, which concludes the proof.

For the failure of c-GNNs we use the example of the proof of Theorem. 4, for which P is incoherent, W is invertible, but any c-GNN is constant. \square

C.4 Additive kernel

Proof of Prop. 8. Again we apply Prop. 6. Assume that x, x' are such that $\Psi_{W,P}(x) = \Psi_{W,P}(x')$ for all one-hop c-SGNN. In particular,

$$\int f(u(v(x) + v(y))) dP(y) = \int f(u(v(x') + v(y))) dP(y)$$

for all MLP f . By taking $f = u^{-1}$, we obtain $v(x) = v(x')$, which leads to $x = x'$ by assumption of injectivity and concludes the proof. \square

C.5 Radial kernel

Proof of Prop. 9. Again we apply Prop. 6. Let x, x' such that $\Psi_{W,P}(x) = \Psi_{W,P}(x')$ for all c-SGNNs. We want to prove that: if P is symmetric, then $x = x'$ or $x = -x'$ (i.e. we quotient $[-1, 1]$ by symmetry), and if P is not symmetric, then necessarily $x = x'$.

By assumption $\int f(W(x, y)) dP(y) = \int f(W(x', y)) dP(y)$ for all MLP f .

If P is symmetric. By choosing $f = (\cdot)^2 \circ w^{-1}$, we have

$$0 = \mathbb{E}(x - X)^2 - \mathbb{E}(x' - X)^2 = x^2 + 2x\mathbb{E}X + \mathbb{E}X^2 - (x')^2 - 2x'\mathbb{E}X - \mathbb{E}X^2 = x^2 - (x')^2$$

which is indeed $x' = x$ or $x' = -x$.

If P is not symmetric. By the previous reasoning, we still have $x' = x$ or $x' = -x$, however, we must now show that the case $x' = -x$ is not possible. By contradiction, assume $x' = -x$ (and $x \neq 0$). Denote $M_k = \mathbb{E}_P X^k$ the k th moment of P . By Lemma 2, P is symmetric iff $M_{2k+1} = 0$ for all k . We are going to show that this is the case by recursion: that is true for $k = 0$ by assumption, and if $M_{2\ell+1} = 0$ for all $\ell \leq k - 1$, by taking $f_0(t) = t^{2k+2}$:

$$\begin{aligned} 0 &= \mathbb{E}(x - X)^{2(k+1)} - \mathbb{E}(x + X)^{2(k+1)} = \sum_{\ell=0}^{2(k+1)} \binom{2(k+1)}{\ell} x^\ell (-1)^\ell M_{2(k+1)-\ell} \\ &\quad - \left(\sum_{\ell=0}^{2(k+1)} \binom{2(k+1)}{\ell} x^\ell M_{2(k+1)-\ell} \right) \\ &= \sum_{\ell=1}^{k+1} \binom{2(k+1)}{2\ell-1} x^{2\ell-1} M_{2(k+1-\ell)+1} \\ &= 2(k+1)xM_{2k+1} \end{aligned}$$

and therefore $M_{2k+1} = 0$, and P is symmetric, which is a contradiction. Therefore, necessarily $x' = x$, which completes the proof. \square

Proof of Prop. 10. Again we apply Prop. 6. Let $x, x' \in \mathbb{S}^{d-1}$ such that $\Psi_{W,P}(x) = \Psi_{W,P}(x')$ for all c-SGNNs. We want to prove that $x = x'$. In particular,

$$\int g(w(x^\top y))dP(y) = \int g(w(x'^\top y))dP(y), \quad (23)$$

for all MLP g .

Recall that we have assumed that that P has a density f decomposed as $f(x) = \sum_{k \geq 0} f_k(x) = \sum_{k \geq 0} \sum_{j=1}^{N(d,k)} a_{k,j} Y_{k,j}(x)$, where $Y_{k,j}$ are spherical harmonics.

Let P_k denote the Legendre/Gegenbauer polynomial of degree k , which satisfies the addition formula

$$P_k(x^\top y) = \frac{1}{N(d,k)} \sum_{j=1}^{N(d,k)} Y_{k,j}(x) Y_{k,j}(y).$$

Then, taking $g = N(d,k)P_k \circ w^{-1}$, note that we have

$$\begin{aligned} \int g(w(x^\top y))dP(y) &= N(d,k) \int P_k(x^\top y) f(y) d\tau(y) \\ &= \sum_j Y_{k,j}(x) \langle f, Y_{k,j} \rangle_{L^2(d\tau)} \\ &= f_k(x). \end{aligned}$$

Thus, (23) implies $f_k(x) = f_k(x')$ for all k . By assumption of injectivity of $x \rightarrow [f_k(x)]_k$, necessarily $x = x'$, which concludes the proof. \square

D Additional material

Lemma 1. *Let (\mathcal{X}, d) be a compact metric space, $\mathcal{F} \subset \cup_{d \geq 1} \mathcal{C}(\mathcal{X}, \mathbb{R}^d)$ be a subspace of continuous multivariate functions on \mathcal{X} that is closed by concatenation, that is, $f, f' \in \mathcal{F} \Rightarrow [f, f'] \in \mathcal{F}$. Define $\mathcal{F}_\rho = \{g \circ f \mid f \in \mathcal{F}, g : \mathbb{R}^d \rightarrow \mathbb{R} \text{ is a MLP with non-linearity } \rho\}$, where ρ is not polynomial. If \mathcal{F} separates points, that is, $\forall x \neq x', \exists f \in \mathcal{F}, f(x) \neq f(x')$, then \mathcal{F}_{MLP} is dense in $\mathcal{C}(\mathcal{X}, \mathbb{R})$ for the supremum norm.*

Proof. The proof uses the classical Stone-Weierstrass theorem: an algebra of continuous functions that separates points is dense in the space of continuous functions (for the supremum norm).

The main point is to check that \mathcal{F}_ρ is an algebra. It is closed by linear combination: for all g, g' MLPs, there is a g'' such that $g \circ f + g' \circ f' = g'' \circ [f, f']$ and \mathcal{F} is closed by concatenation. Closure by multiplication is not true in general, however, following [20], this is true when $\rho = \cos$: since $\cos(a) \cos(b) = \frac{1}{2}(\cos(a+b) - \cos(a-b))$, we have: for $g(x) = \sum_i a_i \cos(b_i^\top x + c_i)$ and similarly g' ,

$$\begin{aligned} (g \circ f) \cdot (g' \circ f') &= \sum_{ij} a_i a'_j \cos(b_i^\top f(x) + c_i) \cos((b'_j)^\top f'(x) + c'_j) \\ &= \sum_{ij} a_i a'_j \frac{1}{2} \left(\cos([b_i, b'_j]^\top [f, f'](x) + c_i + c'_j) \right. \\ &\quad \left. - \cos([b_i, -b'_j]^\top [f, f'](x) + c_i - c'_j) \right) \\ &= g'' \circ [f, f'] \end{aligned}$$

for a certain MLP g'' .

Hence \mathcal{F}_{\cos} is an algebra. Moreover, it separates points: for $x \neq x'$, by hypothesis there is a $f \in \mathcal{F}$ such that $f(x) \neq f(x')$, and by the universality theorem applied to MLPs, this is also true for some $g \circ f$.

To conclude the proof, we note that, by the universality theorem of MLPs, \cos itself can be approached by a MLP with any non-polynomial non-linearity ρ , so that \mathcal{F}_ρ is dense in \mathcal{F}_{\cos} . \square

Lemma 2. *A piecewise continuous function p on $[-1, 1]$ is symmetric iff $\int t^{2k+1}p(t)dt = 0$ for all k .*

Proof. Recall that the Legendre polynomials L_k of degree k are such that: a) they form an orthogonal basis of piecewise continuous functions on $[-1, 1]$ for L^2 , b) respect parity $L_k(-t) = (-1)^k L_k(t)$, c) involves only monomials of the same parity t^{k-2p} , $p = 0, \dots, \lfloor \frac{k}{2} \rfloor$.

By considering the decomposition $p = \sum_k (\int L_k p) L_k$, it is immediate that p is symmetric iff $\int L_{2k+1} p = 0$ for all k , which is the same as $\int t^{2k+1} p(t) dt = 0$ for all k . \square

Lemma 3 (Chaining). *Let $(\mathcal{X}, m_{\mathcal{X}})$ be a compact metric space with diameter $D_{\mathcal{X}}$ and covering numbers $\mathcal{N}(\mathcal{X}, m_{\mathcal{X}}, \varepsilon) \propto \varepsilon^{-d_{\mathcal{X}}}$, and \mathcal{Y} a measurable space. Consider a bivariate measurable function $U : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ that is uniformly C_U -bounded, and L_U -Lipschitz in the first variable. Let y_1, \dots, y_n be drawn i.i.d from a distribution P on \mathcal{Y} . Then, with probability at least $1 - \rho$,*

$$\left\| \frac{1}{n} \sum_i \eta(\cdot, y_i) - \int \eta(\cdot, y) dP(y) \right\|_{\infty} \lesssim \frac{L_U \sqrt{d_{\mathcal{X}}} + (L_U D_{\mathcal{X}} + C_U) \sqrt{\log(1/\rho)}}{\sqrt{n}}.$$

Proof. For any $x \in \mathcal{X}$, define

$$Y_x = \frac{1}{n} \sum_i U(x, x_i) - \int U(x, y) dP(y).$$

Since $|Y_x| \leq 2C_U$, for any fixed $x_0 \in \mathcal{X}$, by Hoeffding's inequality we have: with probability at least $1 - \rho$,

$$|Y_{x_0}| \lesssim C_U \sqrt{\frac{\log(1/\rho)}{n}}.$$

Now we have

$$\left\| \frac{1}{n} \sum_i U(\cdot, x_i) - \int U(\cdot, x) dP(x) \right\|_{\infty} = \sup_{x \in \mathcal{X}} |Y_x| \leq \sup_{x, x' \in \mathcal{X}} |Y_x - Y_{x'}| + |Y_{x_0}|.$$

The second term is bounded by the inequality above. For the first term, we are going to use Dudley's inequality "tail bound" version [40, Thm 8.1.6]. We first check the sub-gaussian increments of the process Y_x . The sub-gaussian norm $\|\cdot\|_{\psi_2}$ is defined in [40, Def. 2.5.6]. For any $x, x' \in \mathcal{X}$, we have

$$\begin{aligned} \|Y_x - Y_{x'}\|_{\psi_2} &\lesssim \left\| \frac{1}{n} \sum_i U(x, y_i) - U(x', y_i) \right\|_{\psi_2} \\ &\lesssim \frac{1}{n} \left(\sum_{i=1}^n \|U(x, y_i) - U(x', y_i)\|_{\psi_2}^2 \right)^{\frac{1}{2}} \\ &\lesssim \frac{1}{n} \left(n \|U(x, \cdot) - U(x', \cdot)\|_{\infty}^2 \right)^{\frac{1}{2}} \\ &\leq \frac{L_U}{\sqrt{n}} m_{\mathcal{X}}(x, x') \end{aligned}$$

where we have used, from [40], Lemma 2.6.8 for the first line, Prop. 2.6.1 for the second, Example 2.5.8 for the third, and the Lipschitz property of U for the last.

Now, we apply Dudley's inequality [40, Thm 8.1.6] to obtain that with probability $1 - \rho$,

$$\begin{aligned} \sup_{x, x' \in \mathcal{X}} |Y_x - Y_{x'}| &\lesssim \frac{L_U}{\sqrt{n}} \left(\int_0^1 \sqrt{\log N(\mathcal{X}, d, \varepsilon)} d\varepsilon + D_{\mathcal{X}} \sqrt{\log(1/\rho)} \right) \\ &\lesssim L_U \frac{\sqrt{d} + D_{\mathcal{X}} \sqrt{\log(1/\rho)}}{\sqrt{n}}. \end{aligned}$$

Combining with the decomposition above yields the desired result. \square

Lemma 4 (Variant of Lemma 6 in [22]). *Let $(E, \|\cdot\|_E)$ be a Banach space and $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$ be a separable Hilbert space. Let L, L' be two bounded operators on E , and $S : E \rightarrow \mathcal{H}$ be a linear operator such that $\|S\|_{\mathcal{H} \rightarrow E} \leq 1$. For $1 \leq i \leq d$ and $1 \leq j \leq d'$, let $h_{ij} = \sum_k \beta_{ijk} \lambda^k$ be a collection of analytic filters, with $B_k = (\beta_{ijk})_{ji} \in \mathbb{R}^{d' \times d}$ the matrix of order- k coefficients, with operator norm $\|B_k\|$. Let $x_1, \dots, x_d \in E$ be a collection of points. Then:*

$$\sqrt{\sum_j \left\| S \sum_i h_{ij}(L)x_i \right\|_{\mathcal{H}}^2} \leq \left(\sum_k \|B_k\| \|L^k\| \right) \sqrt{\sum_i \|x_i\|_E^2} \quad (24)$$

and

$$\sqrt{\sum_j \left\| S \sum_i (h_{ij}(L) - h_{ij}(L'))x_i \right\|_{\mathcal{H}}^2} \leq \sum_k \|B_k\| \sqrt{\sum_i \left(\sum_{\ell=0}^{k-1} \|L^\ell\| \|(L - L')(L')^{k-1-\ell} x_i\|_E \right)^2}. \quad (25)$$

Now, if $\|Lx\|_E \leq \|Sx\|_{\mathcal{H}}$ for some Hilbert space \mathcal{H} , then

$$\sqrt{\sum_j \left\| \sum_i h_{ij}(L)x_i \right\|_E^2} \leq \left(\|B_{0,|\cdot}\| + \sum_{k \geq 1} \|B_k\| \|L^{k-1}\| \right) \sqrt{\sum_i \|x_i\|_E^2}. \quad (26)$$

Proof. The results (24) and (25) are directly from Lemma 6 in [22]. The result (26) is obtained by observing that

$$\begin{aligned} \sqrt{\sum_j \left\| \sum_i h_{ij}(L)x_i \right\|_E^2} &= \sqrt{\sum_j \left\| \sum_{ik} \beta_{ijk} L^k x_i \right\|_E^2} \leq \sum_k \sqrt{\sum_j \left\| \sum_i \beta_{ijk} L^k x_i \right\|_E^2} \\ &\leq \sqrt{\sum_j \left\| \sum_i \beta_{ij0} x_i \right\|_E^2} + \sum_{k \geq 1} \sqrt{\sum_j \left\| S \sum_i \beta_{ijk} L^{k-1} x_i \right\|_{\mathcal{H}}^2}. \end{aligned}$$

We apply (24) on the second term and on the first:

$$\sqrt{\sum_j \left\| \sum_i \beta_{ij0} x_i \right\|_E^2} \leq \sqrt{\sum_j \left(\sum_i |\beta_{ij0}| \|x_i\|_E \right)^2} \leq \|B_{0,|\cdot}\| \sqrt{\sum_i \|x_i\|_E^2}.$$

□

Lemma 5 (Properties of c-GNNs). *Apply a c-GNN to a random graph model. Denote by $f^{(\ell)}$ the function at each layer. Then we have*

$$\|f^{(\ell)}\|_* \leq \|f\|_* \prod_{s=0}^{\ell-1} H_*^{(s)} + \sum_{s=0}^{\ell-1} \|b^{(s)}\| \prod_{p=s+1}^{\ell-1} H_*^{(p)} \quad (27)$$

where $*$ indicates $L^2(P)$ or ∞ .

Moreover, for $x, x' \in \mathcal{X}$, we have

$$\begin{aligned} \|f^{(\ell)}(x) - f^{(\ell)}(x')\| &\leq \left(\prod_{s=0}^{\ell-1} \|B_0^{(s)}\| \right) \|f^{(0)}(x) - f^{(0)}(x')\| \\ &+ L_W d_{\mathcal{X}}(x, x') \sum_{s=0}^{\ell-1} \left(\prod_{p=s+1}^{\ell-1} \|B_0^{(p)}\| \right) \left(\|f^{(0)}\|_{L^2(P)} \prod_{p=0}^s H_2^{(p)} + \sum_{p=0}^{s-1} \|b^{(p)}\| \prod_{r=p+1}^s H_2^{(r)} \right). \end{aligned}$$

Proof. For $j \leq d_\ell$, using Lemma 4, the Lipschitzness of ρ and the easy fact that $\|T_{W,P}f\|_\infty \leq \|f\|_{L^2(P)}$, we write

$$\begin{aligned} \|f^{(\ell)}\|_* &\leq \sqrt{\sum_j \left\| \sum_{i=1}^{d_{\ell-1}} h_{ij}^{(\ell-1)}(T_{W,P})f_i^{(\ell-1)} + b_j^{(\ell-1)} \right\|_*^2} \\ &\leq \sqrt{\sum_j \left\| \sum_{i=1}^{d_{\ell-1}} h_{ij}^{(\ell-1)}(T_{W,P})f_i^{(\ell-1)} \right\|_*^2} + \|b^{(\ell-1)}\| \\ &\leq H_*^{(\ell-1)} \|f^{(\ell-1)}\|_* + \|b^{(\ell-1)}\|. \end{aligned}$$

An easy recursion gives the result.

Now,

$$\begin{aligned} \|f^{(\ell)}(x) - f^{(\ell)}(x')\| &\leq \sqrt{\sum_j \left(\sum_{i=1}^{d_{\ell-1}} h_{ij}^{(\ell-1)}(T_{W,P})f_i^{(\ell-1)}(x) - h_{ij}^{(\ell-1)}(T_{W,P})f_i^{(\ell-1)}(x') \right)^2} \\ &= \sum_k \left\| B_k^{(\ell-1)} \left[T_{W,P}^k f_i^{(\ell-1)}(x) - T_{W,P}^k f_i^{(\ell-1)}(x') \right]_{i=1}^{d_{\ell-1}} \right\| \end{aligned}$$

and since $|T_{W,P}f(x) - T_{W,P}f(x')| \leq L_W d_{\mathcal{X}}(x, x') \|f\|_{L^2(P)}$ and $\|T_{W,P}\|_{L^2(P)} \leq 1$ by Schwartz inequality,

$$\|f^{(\ell)}(x) - f^{(\ell)}(x')\| \leq \|B_0^{(\ell-1)}\| \|f^{(\ell-1)}(x) - f^{(\ell-1)}(x')\| + H_2^{(\ell-1)} \|f^{(\ell-1)}\|_{L^2(P)} L_W d_{\mathcal{X}}(x, x').$$

Again we obtain the result by recursion. \square

Lemma 6. (Properties of c -SGNNs) Denote by $f^{(\ell)}$ the bivariate functions propagated in the inner part of a c -SGNN. We have

$$\|f^{(\ell)}\|_\infty \leq C_\eta \prod_{s=0}^{\ell-1} H_\infty^{(s)} + \sum_{s=0}^{\ell-1} \|b^{(s)}\| \prod_{p=s+1}^{\ell-1} H_\infty^{(p)}.$$

Moreover,

$$\|f^{(\ell)}(\cdot, x) - f^{(\ell)}(\cdot, x')\|_\infty \leq L_\eta d_{\mathcal{X}}(x, x') \prod_{s=0}^{\ell-1} H_\infty^{(s)}.$$

Proof. The first inequality is proved exactly as Lemma 5, noting that $\|T_{W,P}\|_\infty$ even for bivariate functions and $\|\eta\|_\infty \leq C_\eta$.

Then we have

$$\begin{aligned} \|f^{(\ell)}(\cdot, x) - f^{(\ell)}(\cdot, x')\|_\infty &\leq \sqrt{\sum_j \left\| \sum_{i=1}^{d_{\ell-1}} h_{ij}^{(\ell-1)}(T_{W,P}) \left[f_i^{(\ell-1)}(\cdot, x) - f_i^{(\ell-1)}(\cdot, x') \right] \right\|_\infty^2} \\ &\leq H_\infty^{(\ell-1)} \|f^{(\ell-1)}(\cdot, x) - f^{(\ell-1)}(\cdot, x')\|_\infty. \end{aligned}$$

\square

Theorem 5 ([25]). Let A be a $n \times n$ symmetric matrix with independent Bernoulli entries $a_{ij} \sim \alpha_n p_{ij}$. Assume that $\alpha_n \gtrsim \frac{\log n}{n}$. Then, for all $\nu > 0$, there is a constant C_ν such that, for all n , with probability at least $1 - n^{-\nu}$:

$$\frac{1}{n} \left\| \frac{A}{\alpha_n} - P \right\| \leq \frac{C_\nu}{\sqrt{\alpha_n n}}. \quad (28)$$

E Details of numerical experiments

Figure 1. In this figure, we consider a 2-communities SBM with incoherent P , invertible W , but constant degree function. We use dense random edges with $\alpha_n = 1$. We train a permutation-equivariant GNN and a two-hop filtering SGNN on 5 random graphs with $n = 80$ nodes, output dimension $d_{out} = K$, with cross-entropy loss and the Adam optimizer. The displayed graph signal corresponds to the first dimension of the log-softmax of the output. The test graph has $n = 300$ nodes. The graph filters have order 1, such that we actually manipulate the message-passing version of GNNs. The GNN has $M = 5$ hidden layers with internal dimension $d_\ell = 250$ (except $d_0 = 1$ and $d_{out} = 2$) and is trained for 2000 epochs. Each of the GNNs constituting the SGNN has $M = 2$ hidden layers with dimension $d_\ell = 50$ and is trained for 1000 epochs.

Figure 2. We compare one-hop and two-hop input filtering for a simple permutation-invariant SGNN, between the deterministic edges case and the random edges case. We know that the deterministic edges case converges to the c-SGNN in all settings, and we test if the random edges case converge to the deterministic one. We average over 50 random graphs with Gaussian kernel and a range of n 's with $\alpha_n \sim n^{-1/3}$, such that Prop. 1 applies in the two-hop case. The dominating term in the theoretical rate is $\mathcal{O}(1/(\alpha_n \sqrt{n}))$ from Prop. 1.

Figure 3. Here we consider $\mathcal{X} = [-1, 1]$ with Gaussian kernel, and either a symmetric P (uniform) or a non-symmetric but centered P (here a well-adjusted affine by part function). We use random edges with $\alpha_n \sim n^{-1/3}$. We train a SGNN with two-hop input filtering to approximate either a symmetric function $x \rightarrow \cos(5x)$ or a non-symmetric one $x \rightarrow \sin(5x)$ with a simple square loss. We use 5 training graphs of size $n = 150$ and display a test graph with size $n = 400$.