

How does over-squashing affect the power of GNNs ?

F. Di Giovanni and K. Rusch and M. Bronstein and A. Deac and

M. Lackenby and S. Mishra and P. Velickovic

Research Report No. 2023-27
June 2023

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich
Switzerland

How does over-squashing affect the power of GNNs?

Francesco Di Giovanni*
University of Cambridge
fd405@cam.ac.uk

T. Konstantin Rusch*
ETH Zürich
konstantin.rusch@sam.math.ethz.ch

Michael M. Bronstein
University of Oxford

Andreea Deac
Mila, Université de Montréal

Marc Lackenby
University of Oxford

Siddhartha Mishra
ETH Zürich

Petar Veličković
Google DeepMind

Abstract

Graph Neural Networks (GNNs) are the state-of-the-art model for machine learning on graph-structured data. The most popular class of GNNs operate by exchanging information between adjacent nodes, and are known as Message Passing Neural Networks (MPNNs). Given their widespread use, understanding the expressive power of MPNNs is a key question. However, existing results typically consider settings with uninformative node features. In this paper, we provide a rigorous analysis to determine which function classes of node features can be learned by an MPNN of a given capacity. We do so by measuring the level of *pairwise interactions* between nodes that MPNNs allow for. This measure provides a novel quantitative characterization of the so-called over-squashing effect, which is observed to occur when a large volume of messages is aggregated into fixed-size vectors. Using our measure, we prove that, to guarantee sufficient communication between pairs of nodes, the capacity of the MPNN must be large enough, depending on properties of the input graph structure, such as commute times. For many relevant scenarios, our analysis results in impossibility statements in practice, showing that *over-squashing hinders the expressive power of MPNNs*. We validate our theoretical findings through extensive controlled experiments and ablation studies.

1 Introduction

Graphs describe the relational structure for a large variety of natural and artificial systems, making learning on graphs imperative in many contexts [54, 22, 57]. To this end, Graph Neural Networks (GNNs) [29, 51] have emerged as a widely popular framework for graph machine learning, with plentiful success stories in science [5, 18, 9, 52, 36] and technology [42, 41, 21].

Given an underlying graph and features, defined on its nodes (and edges), as inputs, a GNN learns parametric functions from data. Due to the ubiquity of GNNs, characterizing their **expressive power**, i.e., which class of functions a GNN is able to learn, is a problem of great interest. In this context, most available results in literature on the universality of GNNs pertain to impractical higher-order tensors [39, 33] or unique node identifiers that may break the symmetries of the problem [37]. In particular, these results do not necessarily apply to Message Passing Neural Networks (MPNNs) [28], which have emerged as the most popular class of GNN models in recent years. Concerning expressivity results for MPNNs, the most general available characterization is due to [59] and [43], who proved that MPNNs are, at most, as powerful as the Weisfeiler-Leman graph isomorphism test [56] in distinguishing graphs *without any features*. This brings us to an important question:

* Equal contribution.




Figure 1: We study the power of MPNNs in terms of the mixing they induce among features and show that this is affected by the model (via norm of the weights and depth) and the graph topology (via commute times). For the given graph, the MPNN learns stronger mixing (tight springs) for nodes v, u and u, w since their commute time is small, while nodes u, q and u, z , with high commute-time, have weak mixing (loose springs). We characterize over-squashing as the inverse of the mixing induced by an MPNN and hence relate it to its power. In fact, the MPNN might require an impractical depth to solve tasks on the given graph that depend on high-mixing of features assigned to u, z .

Which classes of functions can MPNNs of a given capacity learn, *if node features are specified?*

Razin et al. [46] address this question by characterizing the separation rank of MPNNs; however, their analysis only covers *unconventional* architectures that do not correspond to MPNN models used in practice. In contrast, Alon and Yahav [2] investigate this question *empirically*, by observing that MPNNs fail to solve tasks which involve *long-range interactions* among nodes. This limitation was ascribed to a phenomenon termed as **over-squashing**, which loosely entails messages being ‘squashed’ into fixed-size vectors when the receptive field of a node grows too fast. This effect was formalized in [53, 24, 8], who showed that the Jacobian of the nodes features is affected by topological properties of the graph, such as curvature and effective resistance. However, all the aforementioned papers ignore the specifics of the *task* at hand, i.e., the underlying *function* that the MPNN seeks to learn, leading us to the following question:

How does *over-squashing* affect the expressive power of MPNNs? Can we *measure* it?

Contributions. Our main goal is to address the above question and relate over-squashing to the expressive power of MPNNs. We start by measuring the extent to which an MPNN allows any pair of nodes in the graph to *interact* (via **mixing** their features). With this measure as a tool, we characterize which functions of node features can be learned by an MPNN and how the model architecture and parameters, as well as the topology of the graph, affect the expressive power. More concretely,

- We derive upper bounds on the Hessian of graph-level functions learned by MPNNs and show how their ability to mix features (i.e., model interactions) depends on *both* its *architecture* – via norm of the weights and depth – and the graph *topology* – via commute times.
- We characterize over-squashing as the reciprocal of the maximal mixing that MPNNs can induce among nodes: *the higher this measure, the smaller the class of functions MPNNs can learn.*
- We prove that the weights and depth must be sufficiently large – depending on the topology – to ensure mixing. For some tasks, *the depth must exceed the highest commute time on the graph*, resulting in *impossibility* statements. Our results show that MPNNs fail to learn functions with strong mixing among nodes at high commute time, even when meaningful features are provided.
- We illustrate our theoretical results with controlled experiments that verify our analysis, by highlighting the impact of the architecture (depth), of the topology (commute time), and of the underlying task (the level of mixing required).

Thus, we present a novel paradigm for deriving rigorous results in relating two key notions in graph machine learning, that of GNN expressive power and over-squashing.

2 The Message-Passing paradigm

Definitions on graphs. We denote a graph by $G = (V, E)$, where V is the set of n nodes while E are the edges. We assume that G is *undirected*, *connected* and non-bipartite and define the $n \times n$ adjacency matrix \mathbf{A} as $A_{vu} = 1$ if $(v, u) \in E$ and zero otherwise. We let \mathbf{D} be the diagonal degree matrix with $D_{vv} := d_v$ and use d_{\max} and d_{\min} to denote the maximal and minimal degrees. Since we are interested in the over-squashing phenomenon, which affects the propagation of information, we need to quantify distances on G . We let $d_G(v, u)$ be the length of the shortest path connecting nodes v and u (geodesic distance). While d_G describes how far two nodes u, v are in G , it does not account for how many different routes they can use to communicate. In fact, we will see below that the over-squashing of nodes v, u and, more generally, the mixing induced by MPNNs among the features associated with v, u , can be better quantified by their *commute time* $\tau(v, u)$, equal to the expected number of steps for a random walk to start at v , reach u , and then come back to v .

The MPNN-class. For most problems, graphs are equipped with features $\{\mathbf{x}_v\}_{v \in V} \subset \mathbb{R}^d$, whose matrix representation is $\mathbf{X} \in \mathbb{R}^{n \times d}$. To study the interactions induced by a GNN among pairs of features, we focus on *graph-level* tasks – in Section E of the Appendix, we extend the discussion and our main theoretical results to *node-level* tasks. The goal then is to predict a function $\mathbf{X} \mapsto y_G(\mathbf{X})$, where we assume that the graph G is fixed and thus $y_G : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ is a function of the node features. MPNNs define a family of parametric functions through iterative local updates of the node features: the feature of node v at layer t is derived as

$$\mathbf{h}_v^{(t)} = f^{(t)}\left(\mathbf{h}_v^{(t-1)}, g^{(t)}\left(\{\{\mathbf{h}_u^{(t-1)}, (v, u) \in E\}\}\right)\right), \quad \mathbf{h}_v^{(0)} = \mathbf{x}_v \quad (1)$$

where $f^{(t)}, g^{(t)}$ are learnable functions and the aggregation function $g^{(t)}$ is invariant to permutations. Specifically, we study a class of MPNNs of the following form,

$$\mathbf{h}_v^{(t)} = \sigma\left(\mathbf{\Omega}^{(t)}\mathbf{h}_v^{(t-1)} + \mathbf{W}^{(t)} \sum_u A_{vu} \psi^{(t)}(\mathbf{h}_v^{(t-1)}, \mathbf{h}_u^{(t-1)})\right), \quad \mathbf{h}_v^{(0)} = \mathbf{x}_v, \quad (2)$$

where σ acts pointwise, $\mathbf{\Omega}^{(t)}, \mathbf{W}^{(t)} \in \mathbb{R}^{d \times d}$ are weight matrices, $\mathbf{A} \in \mathbb{R}^{n \times n}$ is *any* matrix satisfying $A_{vu} > 0$ if $(v, u) \in E$ and zero otherwise – \mathbf{A} is typically some (normalized) version of the adjacency matrix \mathbf{A} – and $\psi^{(t)}$ is a learnable message function. The layer-update in (2) includes several common MPNN-models such as GCN [34], SAGE [30], GIN [59], and GatedGCN [10]. For graph-level tasks, an additional permutation-invariant readout READ is required – usually MAX, MEAN, or SUM. We finally define the graph-level function computed by the MPNN after m layers to be

$$y_G^{(m)}(\mathbf{X}) = \langle \boldsymbol{\theta}, \text{READ}(\{\{\mathbf{h}_v^{(m)}\}\}) \rangle, \quad (3)$$

for some learnable $\boldsymbol{\theta} \in \mathbb{R}^d$. While a more complex non-linear readout can be considered, we restrict to a single linear-layer since we are interested in the mixing among variables induced by the MPNN itself through the topology (and not in readout, independently of the graph-structure).

3 On the mixing induced by Message Passing Neural Networks

As one of the main contributions of this paper, we propose a new framework for characterizing the expressive power of MPNNs by *estimating the amount of mixing they induce among pairs of features* \mathbf{x}_v and \mathbf{x}_u , corresponding to the nodes v and $u \in V$. We fix the underlying graph G and denote the ground-truth function to be learned as y_G .

Definition 3.1. For a (smooth) function f of n variables $\{\mathbf{x}_i \in \mathcal{D}_i \subset \mathbb{R}^d\}$, the (maximal) **mixing** induced by f among variables \mathbf{x}_v and \mathbf{x}_u , with $v \neq u$, is

$$\text{mix}_f(v, u) = \max_{\mathbf{x}_i \in \mathcal{D}_i} \max_{1 \leq \alpha, \beta \leq d} \left| \frac{\partial^2 f(\mathbf{X})}{\partial x_v^\alpha \partial x_u^\beta} \right|. \quad (4)$$

Let $f = y_G$ in (4), then the mixing is a measure of the interaction among the features associated with nodes v, u with respect to the underlying task. For example, if $f(\mathbf{X}) = \phi_0(\mathbf{x}_v) + \phi_1(\mathbf{x}_u)$, then the mixing would be zero independently of ϕ_i , since the variables are fully separable. Conversely, if $f(\mathbf{X}) = \phi(\langle \mathbf{x}_v, \mathbf{x}_u \rangle)$, then the amplitude of mixing depends on how *nonlinear* ϕ is. While our approach resonates with the idea of studying the separation rank of an MPNN [46], in contrast to [46], whose analysis applies to a single MPNN with unconventional product aggregations, our results hold for any MPNN of the form (2), which encompasses the majority of architectures used in practice.

Problem statement. Our goal is to study the expressive power of MPNNs in terms of the mixing they can generate among nodes v, u . A low value of mixing implies that the MPNN cannot learn functions that require high mixing of the features associated with v, u . We investigate how weights and depth on the one side, and the graph topology on the other, affect the mixing of an MPNN.

3.1 Pairwise mixing induced by MPNNs

In this Section, our goal is to derive an upper bound on the maximal mixing induced by MPNNs, as defined above, over the features associated with pairs of nodes v, u . To *motivate the structure of this bound*, we consider the simple yet illustrative setting of a linear MPNN in (2) with scalar features, weights $\omega, \mathbf{w} > 0$ and a linear message function of the form $\psi(x, y) = c_1x + c_2y$, for some learnable constants c_1, c_2 . In this case, the layer-update (2) takes the very simple form,

$$h_v^{(t)} = \mathbf{w}(\mathbf{S}\mathbf{h}^{(t-1)})_v, \quad \mathbf{S} := \frac{\omega}{\mathbf{w}}\mathbf{I} + c_1\text{diag}(\mathbf{A}\mathbf{1}) + c_2\mathbf{A} \in \mathbb{R}^{n \times n}, \quad (5)$$

where $\mathbf{1} \in \mathbb{R}^n$ is the vector of ones. Hence, the operator $\mathbf{w}\mathbf{S}$ governs the flow of information from layer $t - 1$ to layer t . Then, it is straightforward to observe that the k -power of this matrix $(\mathbf{w}\mathbf{S})^k$ determines the propagation of information on the graph over k layers, i.e. over walks of length k .

Suppose that nodes v, u are at distance r on the graph. If this *linear* MPNN has depth m , to have any mixing among the features of v and u , there has to be a node i that receives information from both v and u , which only occurs when $m \geq r/2$. If $m = r/2$, then we can estimate the mixing at node i as $((\mathbf{w}\mathbf{S})^m)_{iv}((\mathbf{w}\mathbf{S})^m)_{iu}$. When $m > r/2$, information from v and u can arrive at i after being processed at an intermediate node j ; in fact, this node j first collects information from both v and u over a walk of length $m - k$, and then shares it with node i over a walk of length k . We then obtain terms of the form $\sum_k \sum_j ((\mathbf{w}\mathbf{S})^{m-k})_{jv}((\mathbf{w}\mathbf{S})^k)_{ij}((\mathbf{w}\mathbf{S})^{m-k})_{ju}$, since we are *summing over all possible ways of aggregating information from v and u across intermediate nodes j , before mixing it at i* .

A similar argument to upper bound the maximal mixing also works in the general case of (2), once we account for bounds on the non-linear activation function σ by $c_\sigma = \max\{|\sigma'|, |\sigma''|\}$, and we choose $\omega, \mathbf{w}, c_1, c_2$ satisfying

$$\|\Omega^{(t)}\| \leq \omega, \quad \|\mathbf{W}^{(t)}\| \leq \mathbf{w}, \quad \|\nabla_i \psi^{(t)}\| \leq c_i,$$

for $i = 1, 2$, where $\nabla_i \psi$ is the Jacobian of ψ with respect to the i -th variable, and $\|\cdot\|$ is the *spectral norm* of a matrix. When we consider models with linear message functions, such as GCN, SAGE or GIN, these constants will suffice in determining the upper bound on the mixing of the MPNN. However, in the more general case of non-linear message functions ψ , which for example includes GatedGCN, we also need to account for the term \mathbf{Q}_k , defined below, which arises due to the Hessian of ψ when taking second-order derivatives of the MPNN (2):

given \mathbf{S} in (5), we set

$$\begin{aligned}\mathbf{P}_k &:= (\mathbf{S}^{m-k-1})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k) (\mathbf{A} \mathbf{S}^{m-k-1}) \\ \mathbf{Q}_k &:= \mathbf{P}_k + \mathbf{P}_k^\top + (\mathbf{S}^{m-k-1})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k (\text{diag}(\mathbf{A}\mathbf{1}) + \mathbf{A})) \mathbf{S}^{m-k-1}.\end{aligned}\quad (6)$$

We assume that the Hessian of ψ is bounded as $\|\nabla^2 \psi^{(t)}\| \leq c^{(2)}$. Recall that $y_{\mathbf{G}}^{(m)}$ is the MPNN-prediction (3) and that $\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u)$ is its maximal mixing of nodes v, u as per Definition 3.1.

Theorem 3.2. *Consider an MPNN of depth m as in (2), where σ and $\psi^{(t)}$ are \mathcal{C}^2 functions and we denote the bounds on their derivatives and on the norm of the weights as above. Let \mathbf{S} and \mathbf{Q}_k be defined as in (5) and (6), respectively. If the readout is MAX, MEAN or SUM and $\boldsymbol{\theta}$ in (3) has unit norm, then the mixing $\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u)$ induced by the MPNN over the features of nodes v, u satisfies*

$$\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u) \leq \sum_{k=0}^{m-1} (c_\sigma \mathbf{w})^{2m-k-1} \left(\mathbf{w} (\mathbf{S}^{m-k})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k) \mathbf{S}^{m-k} + c^{(2)} \mathbf{Q}_k \right)_{vu}. \quad (7)$$

Theorem 3.2 shows how the mixing induced by an MPNN depends on the model (via regularity of σ , norm of the weights \mathbf{w} , and depth m) and on the graph-topology (via the powers of \mathbf{A} , which enters the definition of \mathbf{S} in (5)). Despite its generality, from (7) it is as yet unclear which specific properties of the underlying graph affect the mixing the most, and how weights and depth of the MPNN can compensate for it. Therefore, our goal now is to expand (7) and relate it to known quantities on the graph and show how this can be used to characterize the phenomenon of over-squashing. First, we introduce a notion of *capacity* of an MPNN in the spirit of [37].

The capacity of an MPNN. For simplicity, we **assume** that $c_\sigma = 1$, since this is satisfied by most commonly used non-linear activations – it is straightforward to extend the analysis to arbitrary c_σ .

Definition 3.3. *Given an MPNN with m layers and \mathbf{w} the maximal spectral norm of the weights, we say that the pair (m, \mathbf{w}) represents the **capacity** of the MPNN.*

A larger capacity, by increasing either m or \mathbf{w} , or both, heuristically implies that the MPNN has more expressive power and can hence induce larger mixing among the features associated with nodes v, u . Accordingly, given v, u , we formulate the problem of expressivity as: *what is the capacity required to induce enough mixing $\text{mix}_{y_{\mathbf{G}}}(v, u)$?*

4 Over-squashing limits the expressive power of MPNNs

Over-squashing was originally described in [2] as the failure of MPNNs to propagate information across distant nodes. In fact, [8, 24] showed that over-squashing – interpreted as the sensitivity of node v to the input feature at node u via their Jacobian – is affected by the effective resistance (or equivalently, commute time) of v and u . Over-squashing is then related to the inability of MPNNs to model interactions among certain nodes, depending on the underlying graph topology. In light of Theorem 3.2, we propose a *much broader* characterization of over-squashing in terms of the maximal mixing that an MPNN of capacity (m, \mathbf{w}) can induce over nodes v, u on a given graph \mathbf{G} :

Definition 4.1. *Given an MPNN with capacity (m, \mathbf{w}) , we define the **over-squashing** of v, u as*

$$\text{OSQ}_{v,u}(m, \mathbf{w}) := \left(\sum_{k=0}^{m-1} \mathbf{w}^{2m-k-1} \left(\mathbf{w} (\mathbf{S}^{m-k})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k) \mathbf{S}^{m-k} + c^{(2)} \mathbf{Q}_k \right)_{vu} \right)^{-1}.$$

First, we note that our measure of (pairwise) over-squashing is a function of both the graph topology (via \mathbf{S} and \mathbf{Q}_k) and the capacity of the model (via m and \mathbf{w}). In the extreme case of $\mathbf{w} = \infty$ (infinite capacity),

we have $\text{OSQ}_{v,u}(m, \infty) = 0$. Conversely, if the network has no bandwidth through the weights ($\mathbf{w} = 0$), then $\text{OSQ}_{v,u}(m, 0) = \infty$. Second, the proposed measure captures the special case of *under-reaching* [6], since $\text{OSQ}_{v,u}$ is infinite (i.e., zero mixing) whenever $2m < d_G(v, u)$, where d_G is the shortest-walk distance on G . Third, we recall that for simplicity we have taken $c_\sigma = 1$, but the measure naturally extends to arbitrary non-linear activations σ . Finally, we generalize the characterization of OSQ to node-level tasks in Section E of the Appendix.

Definition 4.1 measures the (maximal) pairwise mixing induced by an MPNN of capacity (m, \mathbf{w}) : *the higher $\text{OSQ}_{v,u}$, the lower the mixing*. In fact, by Theorem 3.2 we find that a **necessary condition for an MPNN to learn a function y_G with mixing $\text{mix}_{y_G}(v, u)$ is**

$$\text{OSQ}_{v,u}(m, \mathbf{w}) < (\text{mix}_{y_G}(v, u))^{-1}. \quad (8)$$

Therefore, an MPNN of given capacity on a featured graph (G, \mathbf{X}) , might suffer or not from over-squashing, depending on the level of mixing required by the underlying task. We study how OSQ varies with v, u and show that it will be larger for nodes v, u that are far apart, or more generally, have a high commute time, validating that over-squashing is a *phenomenon that strongly depends on the graph-topology*. For a given graph G , to reduce the value of OSQ and hence satisfy (8), the capacity (m, \mathbf{w}) must satisfy constraints posed by G and the choice of v, u . Since we can increase the capacity by taking either larger weights or more layers, we consider these two regimes separately.

Choice of \mathbf{A} . For simplicity, we restrict our analysis to the case $\mathbf{A} = \mathbf{A}_{\text{sym}} := \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ and extend the results to $\mathbf{D}^{-1} \mathbf{A}$ and \mathbf{A} in Section D of the Appendix. For the unnormalized adjacency, the bounds look more favourable, simply because we are no longer normalizing the messages, which makes the model more sensitive for all nodes. In Section D.4 we comment on how one could account for the lack of normalization by taking relative measurement for OSQ in the spirit of [59].

4.1 The case of fixed depth m and variable weights norm \mathbf{w}

To assess the ability of the norm of the weights \mathbf{w} to increase the capacity of an MPNN and hence reduce OSQ, we consider the limit case where the depth m is the minimal required for an MPNN to have a non-zero mixing among v, u (half the shortest-walk distance d_G between the nodes).

Theorem 4.2. *Let $\mathbf{A} = \mathbf{A}_{\text{sym}}$, $r := d_G(v, u)$, $m = \lceil r/2 \rceil$, and q be the number of paths of length r between v and u . For an MPNN satisfying Theorem 3.2 with capacity $(m = \lceil r/2 \rceil, \mathbf{w})$, we find $\text{OSQ}_{v,u}(m, \mathbf{w}) \cdot (c_2 \mathbf{w})^r (\mathbf{A}^T)_{vu} \geq 1$. In particular, if the MPNN generates mixing $\text{mix}_{y_G}(v, u)$, then*

$$\mathbf{w} \geq \frac{d_{\min}}{c_2} \left(\frac{\text{mix}_{y_G}(v, u)}{q} \right)^{\frac{1}{r}}.$$

Theorem 4.2 highlights that if the depth is set as the minimum required for *any* non-zero mixing, then the norm of the weights \mathbf{w} has to be large enough depending on the connectivity of G – recall that for models as GCN, we have $c_2 = 1$. However, increasing \mathbf{w} is not optimal, since it makes the whole model more sensitive for all nodes, which may lead to poorer generalization capabilities [7, 26].

Some examples. We illustrate the bounds in Theorem 4.2 and for simplicity, we set $c_2 = 1$. Consider a *tree* T_d of arity d , with v the root and u a leaf at distance r and depth $m = r/2$; then $\text{OSQ}_{v,u}(m, \mathbf{w}) \geq \mathbf{w}^{-r} (d+1)^{r-1}$ and the spectral norm required to generate mixing $y(v, u)$ is

$$\mathbf{w} \geq (d+1) \left(\frac{y(v, u)}{d+1} \right)^{\frac{1}{r}}.$$

We note that by taking $d = 1$ we recover the case of the path-graph (1D grid). Since the spectral norm of the weights grows with the branching factor, we see that, in general, the capacity required by MPNNs to solve

long-range tasks could be higher on graphs than on sequences [2]. We also consider the case of a 1-layer MPNN on a complete graph K_n with $v \neq u$. We find that $\text{OSQ}_{v,u}(m, \mathbf{w}) \geq (n-1)/w$ and hence the spectral norm required to generate mixing $y(v, u)$ is $w \geq (n-1)y(v, u)$. We note how the measure of over-squashing also captures the problem of redundancy of messages [15]. In fact, even if v, u are at distance 1, the more nodes are there in the complete graph and hence the more messages are exchanged, the more difficult for a shallow MPNN to induce enough mixing among those *specific* nodes.

4.2 The case of fixed weights norm w and variable depth m

We now study the (desirable) setting where w is bounded, and derive the depth necessary to induce mixing of nodes v, u . Below, we let $\Delta = \mathbf{I} - \mathbf{A}_{\text{sym}}$ be the normalized graph Laplacian, whose eigenvalues are $0 = \lambda_0 < \lambda_1 \leq \dots \leq \lambda_{n-1}$; we note that λ_1 is the spectral gap and $\lambda_{n-1} < 2$ if G is not bipartite [17]. We also recall that d_G is the shortest-walk distance and τ is the *commute time* (defined in Section 2). Finally, if d_{\max} and d_{\min} denote the maximal and minimal degrees, respectively, we set $\gamma := \sqrt{d_{\max}/d_{\min}}$.

Theorem 4.3. *Consider an MPNN satisfying Theorem 3.2, with $\max\{w, \omega/w + c_1\gamma + c_2\} \leq 1$, and $\mathbf{A} = \mathbf{A}_{\text{sym}}$. If $\text{OSQ}_{v,u}(m, \mathbf{w}) \cdot (\text{mix}_{y_G}(v, u)) \leq 1$, and hence the MPNN generates mixing $\text{mix}_{y_G}(v, u)$ among the features associated with nodes v, u , then the number of layers m satisfies*

$$m \geq \frac{\tau(v, u)}{4c_2} + \frac{|E|}{\sqrt{d_v d_u}} \left(\frac{\text{mix}_{y_G}(v, u)}{\gamma \mu} - \frac{1}{c_2} \left(\frac{\gamma + |1 - c_2 \lambda^*|^{r-1}}{\lambda_1} + 2 \frac{c^{(2)}}{\mu} \right) \right),$$

where $r = d_G(v, u)$, $\mu = 1 + 2c^{(2)}(1 + \gamma)$ and $|1 - c_2 \lambda^*| = \max_{0 < \ell \leq n-1} |1 - c_2 \lambda_\ell| < 1$.

Theorem 4.3 provides a *necessary condition* on the depth of an MPNN to induce enough mixing among nodes v, u . We see that the MPNN must be sufficiently deep if the task depends on interactions between nodes at high commute time τ and in fact, if the required mixing is large enough, then $m = \Omega(\tau(v, u))$. Note that the lower bound on the depth can translate into a *practical impossibility statement*, since the commute time τ can be as large as $\mathcal{O}(n^3)$ [14]. We also recall that for MPNN models such as GCN, the bounds above simplify since $c^{(2)} = 0$ and hence $\mu = 1$.

Therefore, if (i) the graph is such that the commute time between v, u is large, and (ii) the task depends on high-mixing of features associated with v, u , then *over-squashing limits the expressive power* of MPNNs since the depth has to scale impractically with the graph size n – which also brings forth other issues such as over-smoothing [48, 13, 49, 23]. In contrast to existing analysis on the limitations of MPNNs via graph-isomorphism test, our results characterize the expressivity of MPNNs even when **meaningful features are provided**. In fact, Theorem 4.3 implies that:

Corollary 4.4. *On a graph with features, MPNNs as in Theorem 4.3 with depth $m \leq n$, cannot learn functions that induce high mixing among features of nodes with large commute time.*

We also note that for two adjacent nodes v, u , the commute time equals $2|E|$ if (v, u) is a **cut-edge** [1]. Our result shows that MPNNs may require $m = \Omega(|E|)$ to generate enough mixing of features associated with a cut-edge, drawing a connection with [61], where it was shown that most GNNs fail to identify cut-edges on *unattributed* graphs. Finally, since the commute time is proportional to the effective resistance R by $2R(v, u)|E| = \tau(v, u)$, Theorem 4.3 connects recent works [24, 8] relating the Jacobian of node features and R , and the expressive power of MPNNs in terms of their mixing.

5 Experimental validation of the theoretical results

Next, we aim to empirically verify the theoretical findings of this paper, i.e., the characterization of the impact of the graph topology (via commute time τ), the GNN architecture (depth, norm of weights), and the underlying

task (node mixing) on over-squashing. This, however, requires detailed information about the underlying function to be learned, which is not readily available in practice. Hence, we perform our empirical test in a controlled environment, but at the same time, we base our experiments on the *real world* ZINC chemical dataset [31] and follow the experimental setup in [25], constraining the number of molecular graphs to 12K. Moreover, we exclude the edge features from this experiment and fix the MPNN size to $\sim 100\text{K}$ parameters. However, instead of regressing the constrained solubility based on the molecular input graphs, we define our own synthetic node features as well as our own target values as follows.

Let $\{\mathbf{G}^i\}$ be the set of the 12K ZINC molecular graphs. We set all node features to zero, except for two, which are set to uniform random numbers $x_{u^i}^i, x_{v^i}^i$ between 0 and 1 (i.e., $x_{u^i}^i, x_{v^i}^i \sim \mathcal{U}(0, 1)$) for all i . The target output is set to $y^i = \tanh(x_{u^i}^i + x_{v^i}^i)$ for all i . Hence, the task entails a non-linear mixing with non-vanishing second derivatives. The two non-zero node features $x_{u^i}^i, x_{v^i}^i$ are positioned on \mathbf{G}_i according to the commute time τ , i.e., for a given $\alpha \in [0, 1]$, we choose the nodes u^i, v^i as the α -quantile of the τ -distribution over \mathbf{G}_i . This enables us to have the desired control on the level of commute time of the underlying mixing (i.e., $\alpha_1 < \alpha_2$, if α_1 results in a lower τ than α_2). See Fig. 2 for an illustration of a ZINC graph together with the histogram of the corresponding τ between all pairs of different nodes. We call this graph dataset the *synthetic ZINC dataset*.




Figure 2: **(Left)** Exemplary molecular graph of the ZINC (12K) dataset with colored nodes corresponding to different values of commute time τ . We note that τ is a more refined measure than the distance, and in fact beyond long-range nodes (red case), τ also captures other topological properties (yellow nodes are adjacent but belong to a *cut-edge*, so their commute-time is $2|E|$). **(Right)** Histogram of commute time τ between all pairs of the graph nodes.

We consider four different MPNN models namely GCN [34], GIN [59], GraphSAGE [30], and GatedGCN [10]. Moreover, we choose the MAX-pooling as the GNN readout, which is supported by Theorem 3.2 and forces the GNNs to make use of the message-passing in order to learn the mixing.

5.1 The role of commute time

In this task, we empirically analyse the effect of the commute time τ of the underlying mixing on the performance of the MPNNs. To this end, we fix the architecture for all considered MPNNs. In particular, we set the depth to $m = \max_i \lceil \text{diam}(\mathbf{G}^i)/2 \rceil$, which happens to be $m = 11$ for the considered ZINC 12K graphs, such that the MPNNs are guaranteed *not to underreach*. We further vary the value of the α -quantile of the τ -distributions over the graphs \mathbf{G}^i between 0 and 1, thus controlling the level of commute times. According to our theoretical findings in Section 4, the over-squashing measure $\text{OSQ}_{v,u}$ (Definition 4.1) heavily depends on the commute time τ of the underlying mixing as derived in Theorem 4.3. In fact, this can be empirically verified in Appendix Fig. 7, where $\text{OSQ}_{v,u}$ increases for increasing values of α . Thus, we would expect the MPNNs to perform significantly worse for increasing levels of the commute time. This is indeed confirmed in Fig. 3 depicting the resulting test mean absolute error (MAE), and showing that the test MAE increases for larger values of α for all considered MPNNs. The same qualitative behavior can be observed for the *training* MAE (Appendix F.3).

We further note that the models exhibit low standard deviation (over several trained MPNNs), as can be seen in Appendix Table 2.




Figure 3: Test MAE of GCN, GIN, GraphSAGE, and GatedGCN on synthetic ZINC, where the commute time of the underlying mixing is varied, while the MPNN architecture is fixed (e.g., depth, number of parameters), i.e., mixing according to increasing values of the α -quantile of the τ -distribution over the ZINC graphs.




Figure 4: Test MAE of GCN, GIN, GraphSAGE, and GatedGCN on synthetic ZINC, where the commute time is fixed to be high (i.e., at the level of the 0.8-quantile), while only the depth of the underlying MPNN is varied between 5 and 32 (all other architectural components are fixed).

5.2 The role of depth

In this task, we study the effect of the depth on the performance of the MPNNs. To this end, we consider a high commute time-regime by setting $\alpha = 0.8$. Note that in this case the maximum (over all graphs G^i) shortest path between two nodes u^i, v^i is 14. Therefore, a depth of $m = 7$ is sufficient to avoid under-reaching on all graphs. However, according to the over-squashing measure we provide and the conclusions of Theorem 4.3, we expect the MPNNs to improve with increasing number of layers. This expectation is further evidenced in Appendix Fig. 8, where the computed OSQ decreases for increasing number of layers. In Fig. 4, we plot the test MAE of all considered MPNNs for increasing number of layers. We can indeed see that all considered GNNs benefit from depth, and thus higher capacity (Definition 3.3), as GatedGCN obtains the lowest test MAE with 17 layers, GraphSAGE with 30 layers, GIN with 24 layers, and GCN with 32 layers. The same qualitative behavior can be observed for the *training* MAE (Appendix F.3). We note that we cannot increase the depth arbitrarily, as at some point other issues emerge, such as vanishing gradients, which impact the trainability of the MPNNs [47]. Moreover, as can be seen in Table 3, all models exhibit low standard deviation (over several trained MPNNs). We further note that although a depth of $m = 7$ is sufficient to learn the underlying task, all MPNNs greatly improve their performance by using deeper models.

5.3 The role of mixing

We further test the considered MPNN architectures on their performance with respect to different mixings. To this end, we consider again the tanh-based mixing as in our previous tasks (i.e., regressing targets $y_i = \tanh(x_{u^i}^i + x_{v^i}^i)$ for each graph G^i in the dataset), as well as another mixing based on the exponential function (i.e., with targets $y_i = \exp(x_{u^i}^i + x_{v^i}^i)$). We note that these two tasks differ significantly in terms of their maximal mixing values (4) (shown in Table 1). Thus, according to (8) and Theorem 4.3, we would expect that MPNNs perform significantly worse in the case of higher maximal mixing, i.e., for the exponential-based mixing compared to the tanh-mixing. To confirm this empirically, we train the MPNNs on both types of mixing and provide the resulting relative MAEs (i.e., MAE divided by the L^1 -norm of the targets) in Table 1. We can

see that all four MPNNs perform significantly better on the tanh-mixing than on the exponential-based mixing. Moreover, increasing the range for the exponential-based mixing from 1 to 1.5 further impairs the performance of all considered MPNNs. In order to check if this difference in performance can simply be explained by a higher capacity required by a neural network to accurately approximate the mapping $\exp(x + y)$ compared to $\tanh(x + y)$ for some inputs $x, y \in \mathbb{R}$, we train a simple two-layer feed-forward neural network (with 2 inputs, i.e., x and y) on both mappings. The trained networks reach a similarly low relative MAE of 4.6×10^{-4} for the $\tanh(x + y)$ mapping as well as 4.1×10^{-4} for the $\exp(x + y)$ mapping using an input range of $(0, 1)$ and 4.0×10^{-4} for an input range of $(0, 1.5)$. Thus, we can conclude that the significant differences in the obtained results in Table 1 are not caused by a higher capacity required by a neural network to learn the underlying mappings of the different mixings.

Table 1: Relative MAE of GCN, GIN, GraphSAGE and GatedGCN on different choices of mixing on synthetic ZINC for a fixed 0.8-quantile of the commute time distributions over graphs G_i .

Mixing	input interval	maximal mixing	GCN	GIN	GraphSAGE	GatedGCN
$\tanh(x_{u^i}^i + x_{v^i}^i)$	$(0, 1)$	≈ 0.77	0.024	0.014	0.006	0.004
$\exp(x_{u^i}^i + x_{v^i}^i)$	$(0, 1)$	≈ 7.4	0.043	0.021	0.033	0.008
$\exp(x_{u^i}^i + x_{v^i}^i)$	$(0, 1.5)$	≈ 20.1	0.054	0.035	0.075	0.014

6 Discussion

Related Work. The MPNN class in (2) is as **powerful** as the 1-WL test [56] in distinguishing unattributed graphs [59, 43]. In fact, MPNNs typically struggle to compute graph properties on feature-less graphs [20, 16, 50, 37]. The expressivity of GNNs has also been studied from the point of view of logical and tensor languages [6, 4, 27], which have provided hierarchies similar to color-refinement algorithms. Nonetheless, far less is known about which functions of node features MPNNs can learn and the capacity required to do so. Razin et al. [46] recently studied the separation rank of a specific MPNN class. While this approach is a strong inspiration for our work, the results in [46] only consider the interactions among a subset of the nodes and its complement, and only apply to a single unconventional family of MPNNs which does not include MPNN models used in practice. *Our results instead hold in the full generality of (2) and provide a novel approach to investigate expressivity of MPNNs through the mixing they are able to generate among features.*

The problem of **over-squashing** was introduced in [2] and studied through sensitivity analysis [58] in [53]. This approach was generalized in [8, 24] who proved that the Jacobian of node features is likely to be small if the nodes have high commute time (effective resistance). However, the analysis in [8, 24] does not include all MPNNs as in (2), and, more importantly, it lacks a connection between over-squashing and the expressive power of MPNNs. In fact, considering the first-order Jacobian might not be suited to study over-squashing for graph-level tasks. *We have extended these ideas to connect over-squashing and expressive power by studying higher-order derivatives of the MPNN and relating them to the capacity of the model and the underlying graph-topology.*

6.1 Limitations and ways forward

The measure OSQ. Definition 4.1 considers pairs of nodes and second-order derivatives. This could be generalized to build a hierarchy of measures with increasingly higher-order interactions of nodes. In the paper we focused our analysis on the case of graph-level tasks; we provide an extension to node-level functions in Section E. Finally, we note that OSQ is based on the maximal mixing induced by MPNNs (upper bounds). While this allows us to derive *necessary* conditions on the capacity of the MPNN to learn functions with a given

mixing, even when these criteria are met, the MPNN may still be unable to induce such mixing. Extending the analysis to lower bounds on the mixing, in order to provide *sufficient* conditions for learning tasks, is a topic for future work.

Beyond sum-aggregations. Our results apply to MPNNs as in (2), where \mathbf{A} is constant, and do not include attention-based MPNNs [55, 11] or Graph-Transformers [35, 40, 60, 45] which further depend on features via normalization. Extending the analysis to these models is conceptually similar and only more technically involved. More generally, one could replace the aggregation $\sum_u \mathbf{A}_{vu}$ in (2) with a smooth, permutation invariant operator \oplus [12, 44] and follow our approach to study the mixing abilities of \oplus MPNNs. *Our formalism will then prove useful to assess if different aggregations are more expressive in terms of the mixing (interactions) they are able to generate.*

Graph rewiring. Another way of going beyond (2) to find MPNNs with lower OSQ is to replace \mathbf{A} with a different matrix \mathbf{A}' , (partly) independent of the connectivity of the input graph. In other words, one could rewire the graph, by replacing \mathbf{G} with \mathbf{G}' . Indeed, Theorem 4.3 further validates why recent graph-rewiring methods such as [3, 32, 19, 8] manage to alleviate over-squashing: by adding edges that decrease the overall effective resistance (commute time) of the graph, these methods reduce the measure OSQ. More generally, *Definition 4.1 allows one to measure whether a given rewiring is beneficial in terms of over-squashing (and hence of the mixing generated) and to what extent.*

Acknowledgements

The authors would like to thank Dr. Martín Arjovsky, Dr. Charles Blundell and Dr. Karl Tuyls for their insightful feedback and constructive suggestions on an earlier version of the manuscript.

References

- [1] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *20th Annual Symposium on Foundations of Computer Science (sfcs 1979)*, pages 218–223. IEEE Computer Society, 1979.
- [2] U. Alon and E. Yahav. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*, 2021.
- [3] A. Arnaiz-Rodríguez, A. Begga, F. Escolano, and N. Oliver. DiffWire: Inductive Graph Rewiring via the Lovász Bound. In *The First Learning on Graphs Conference*, 2022. URL <https://openreview.net/pdf?id=IXvfIexOmX6f>.
- [4] W. Azizian and M. Lelarge. Expressive power of invariant and equivariant graph neural networks. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=lxHgXYN4bwl>.
- [5] V. Bapst, T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. S. Schoenholz, A. Obika, A. W. Nelson, T. Back, D. Hassabis, et al. Unveiling the predictive power of static structure in glassy systems. *Nature Physics*, 16(4):448–454, 2020.
- [6] P. Barceló, E. V. Kostylev, M. Monet, J. Pérez, J. Reutter, and J. P. Silva. The logical expressiveness of graph neural networks. In *International Conference on Learning Representations*, 2019.
- [7] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky. Spectrally-normalized margin bounds for neural networks. *Advances in neural information processing systems*, 30, 2017.

- [8] M. Black, A. Nayyeri, Z. Wan, and Y. Wang. Understanding oversquashing in gnns through the lens of effective resistance. *arXiv preprint arXiv:2302.06835*, 2023.
- [9] C. Blundell, L. Buesing, A. Davies, P. Veličković, and G. Williamson. Towards combinatorial invariance for kazhdan-lusztig polynomials. *arXiv preprint arXiv:2111.15161*, 2021.
- [10] X. Bresson and T. Laurent. Residual gated graph convnets. *arXiv preprint arXiv:1711.07553*, 2017.
- [11] S. Brody, U. Alon, and E. Yahav. How attentive are graph attention networks? In *International Conference on Learning Representations*, 2022.
- [12] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- [13] C. Cai and Y. Wang. A note on over-smoothing for graph neural networks. *arXiv preprint arXiv:2006.13318*, 2020.
- [14] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph captures its commute and cover times. *computational complexity*, 6(4):312–340, 1996.
- [15] R. Chen, S. Zhang, Y. Li, et al. Redundancy-free message passing for graph neural networks. *Advances in Neural Information Processing Systems*, 35:4316–4327, 2022.
- [16] Z. Chen, L. Chen, S. Villar, and J. Bruna. Can graph neural networks count substructures? *Advances in neural information processing systems*, 33:10383–10395, 2020.
- [17] F. R. Chung and F. C. Graham. *Spectral graph theory*. American Mathematical Soc., 1997.
- [18] A. Davies, P. Veličković, L. Buesing, S. Blackwell, D. Zheng, N. Tomašev, R. Tanburn, P. Battaglia, C. Blundell, A. Juhász, et al. Advancing mathematics by guiding human intuition with ai. *Nature*, 600(7887):70–74, 2021.
- [19] A. Deac, M. Lackenby, and P. Veličković. Expander graph propagation. In *The First Learning on Graphs Conference*, 2022.
- [20] N. Dehmamy, A.-L. Barabási, and R. Yu. Understanding the representation power of graph neural networks in learning graph topology. *Advances in Neural Information Processing Systems*, 32, 2019.
- [21] A. Derrow-Pinion, J. She, D. Wong, O. Lange, T. Hester, L. Perez, M. Nunkesser, S. Lee, X. Guo, B. Wiltshire, et al. Eta prediction with graph neural networks in google maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 3767–3776, 2021.
- [22] G. DeZoort, P. W. Battaglia, C. Biscarat, and J.-R. Vlimant. Graph neural networks at the large hadron collider. *Nature Reviews Physics*, pages 1–23, 2023.
- [23] F. Di Giovanni, J. Rowbottom, B. P. Chamberlain, T. Markovich, and M. M. Bronstein. Graph neural networks as gradient flows. *arXiv preprint arXiv:2206.10991*, 2022.
- [24] F. Di Giovanni, L. Giusti, F. Barbero, G. Luise, P. Lio, and M. Bronstein. On over-squashing in message passing neural networks: The impact of width, depth, and topology. In *International Conference on Machine Learning*, 2023.
- [25] V. P. Dwivedi, C. K. Joshi, T. Laurent, Y. Bengio, and X. Bresson. Benchmarking graph neural networks. *arXiv:2003.00982*, 2020.
- [26] V. Garg, S. Jegelka, and T. Jaakkola. Generalization and representational limits of graph neural networks. In *International Conference on Machine Learning*, pages 3419–3430. PMLR, 2020.

- [27] F. Geerts and J. L. Reutter. Expressiveness and approximation properties of graph neural networks. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=wIzUeM3TAU>.
- [28] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272. PMLR, 2017.
- [29] M. Gori, G. Monfardini, and F. Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 729–734. IEEE, 2005.
- [30] W. L. Hamilton, R. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1025–1035, 2017.
- [31] J. J. Irwin, T. Sterling, M. M. Mysinger, E. S. Bolstad, and R. G. Coleman. Zinc: a free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7):1757–1768, 2012.
- [32] K. Karhadkar, P. K. Banerjee, and G. Montufar. FoSR: First-order spectral rewiring for addressing oversquashing in GNNs. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=3YjQfCLdrzz>.
- [33] N. Keriven and G. Peyré. Universal invariant and equivariant graph neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [34] T. N. Kipf and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*, 2017.
- [35] D. Kreuzer, D. Beaini, W. Hamilton, V. Létourneau, and P. Tossou. Rethinking graph transformers with spectral attention. In *Advances in Neural Information Processing Systems*, volume 34, pages 21618–21629, 2021.
- [36] G. Liu, D. B. Catacutan, K. Rathod, K. Swanson, W. Jin, J. C. Mohammed, A. Chiappino-Pepe, S. A. Syed, M. Fragis, K. Rachwalski, J. Magolan, M. G. Surette, B. K. Coombes, T. Jaakkola, R. Barzilay, J. J. Collins, and J. M. Stokes. Deep learning-guided discovery of an antibiotic targeting acinetobacter baumannii. *Nature Chemical Biology*, May 2023. ISSN 1552-4469. doi: 10.1038/s41589-023-01349-8. URL <https://doi.org/10.1038/s41589-023-01349-8>.
- [37] A. Loukas. What graph neural networks cannot learn: depth vs width. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=B1l2bp4YwS>.
- [38] L. Lovász. Random walks on graphs. *Combinatorics, Paul erdos is eighty*, 2(1-46):4, 1993.
- [39] H. Maron, E. Fetaya, N. Segol, and Y. Lipman. On the universality of invariant networks. In *International conference on machine learning*, pages 4363–4371. PMLR, 2019.
- [40] G. Mialon, D. Chen, M. Selosse, and J. Mairal. Graphit: Encoding graph structure in transformers. *CoRR*, abs/2106.05667, 2021.
- [41] A. Mirhoseini, A. Goldie, M. Yazgan, J. W. Jiang, E. Songhori, S. Wang, Y.-J. Lee, E. Johnson, O. Pathak, A. Nazi, et al. A graph placement methodology for fast chip design. *Nature*, 594(7862):207–212, 2021.
- [42] F. Monti, F. Frasca, D. Eynard, D. Mannion, and M. M. Bronstein. Fake news detection on social media using geometric deep learning. *arXiv preprint arXiv:1902.06673*, 2019.
- [43] C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. E. Lenssen, G. Rattan, and M. Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *AAAI Conference on Artificial Intelligence*, pages 4602–4609. AAAI Press, 2019.
- [44] E. Ong and P. Veličković. Learnable commutative monoids for graph neural networks. *arXiv preprint arXiv:2212.08541*, 2022.

- [45] L. Rampasek, M. Galkin, V. P. Dwivedi, A. T. Luu, G. Wolf, and D. Beaini. Recipe for a general, powerful, scalable graph transformer. In *Advances in Neural Information Processing Systems*, 2022.
- [46] N. Razin, T. Verbin, and N. Cohen. On the ability of graph neural networks to model interactions between vertices. *arXiv preprint arXiv:2211.16494*, 2022.
- [47] T. K. Rusch, B. Chamberlain, J. Rowbottom, S. Mishra, and M. Bronstein. Graph-coupled oscillator networks. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 18888–18909. PMLR, 2022.
- [48] T. K. Rusch, M. M. Bronstein, and S. Mishra. A survey on oversmoothing in graph neural networks. *arXiv preprint arXiv:2303.10993*, 2023.
- [49] T. K. Rusch, B. P. Chamberlain, M. W. Mahoney, M. M. Bronstein, and S. Mishra. Gradient gating for deep multi-rate learning on graphs. In *International Conference on Learning Representations*, 2023.
- [50] R. Sato, M. Yamada, and H. Kashima. Approximation ratios of graph neural networks for combinatorial problems. *Advances in Neural Information Processing Systems*, 32, 2019.
- [51] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
- [52] J. M. Stokes, K. Yang, K. Swanson, W. Jin, A. Cubillos-Ruiz, N. M. Donghia, C. R. MacNair, S. French, L. A. Carfrae, Z. Bloom-Ackermann, et al. A deep learning approach to antibiotic discovery. *Cell*, 180(4): 688–702, 2020.
- [53] J. Topping, F. Di Giovanni, B. P. Chamberlain, X. Dong, and M. M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. In *International Conference on Learning Representations*, 2022.
- [54] P. Veličković. Everything is connected: Graph neural networks. *Current Opinion in Structural Biology*, 79: 102538, 2023.
- [55] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [56] B. Weisfeiler and A. Leman. The reduction of a graph to canonical form and the algebra which appears therein. *nti, Series*, 2(9):12–16, 1968.
- [57] G. Williamson. Is deep learning a useful tool for the pure mathematician? *arXiv preprint arXiv:2304.12602*, 2023.
- [58] K. Xu, C. Li, Y. Tian, T. Sonobe, K.-i. Kawarabayashi, and S. Jegelka. Representation learning on graphs with jumping knowledge networks. In *International Conference on Machine Learning*, pages 5453–5462. PMLR, 2018.
- [59] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- [60] C. Ying, T. Cai, S. Luo, S. Zheng, G. Ke, D. He, Y. Shen, and T.-Y. Liu. Do transformers really perform badly for graph representation? In *Advances in Neural Information Processing Systems*, volume 34, pages 28877–28888, 2021.
- [61] B. Zhang, S. Luo, L. Wang, and D. He. Rethinking the expressive power of gnns via graph biconnectivity. In *The Eleventh International Conference on Learning Representations*, 2023.

A Outline of the appendix

We provide an overview of the appendix. Since in the appendix we report additional theoretical results and considerations, we first point out to the most relevant content: the proofs of the main results, the extension of our discussion and analysis to node-level tasks, and the additional ablation studies.

Where to find proofs of the main results. We prove Theorem 3.2 in Section C.1, we prove Theorem 4.2 in Section D.1, and finally we prove Theorem 4.3 in Section D.3.

Where to find the extension to node-level tasks. Concerning the case of node-level tasks, we present a thorough discussion on the matter in Section E, where we extend the definition of the over-squashing measure and generalize Theorem 3.2 and Theorem 4.3 to node-level predictions of the MPNN class in (2).

Where to find additional ablation studies. In Section F we have conducted further experiments on the profile of the over-squashing measure OSQ across different MPNN models as well as on the training mean average error, to further validate our claims on over-squashing hindering the expressive power of MPNNs.

Next, we summarize the contents of the Appendix more in detail below.

- In order to be self-consistent, in Section B we review important notions pertaining to the spectrum of the graph-Laplacian and known properties of random walks on graphs, that will be then be used in our proofs.
- In Section C we prove the main theorem on the maximal mixing induced by MPNNs (Theorem 3.2). In particular, we also derive additional results on the mixing generated at a specific node, which will turn out useful when extending the characterization of the over-squashing measure OSQ for node-level tasks.
- In Section D we prove the main results of Section 4, mainly Theorem 4.2 and Theorem 4.3. Further, we also derive an explicit (sharper) characterization of the depth required to induce enough mixing among nodes, in terms of the pseudo-inverse of the graph-Laplacian. Finally, in Section D.4 we extend the results to the case of the unnormalized adjacency matrix and discuss relative over-squashing measures.
- In Section E we generalize the over-squashing measure for node-level tasks, commenting on the differences between our approach and existing works (mainly [53, 8, 24]). In particular, we show that the same conclusions of Theorem 4.3 hold for node-level predictions too.
- Finally, in Section F we report additional details on our experimental setup and further ablation studies concerning the over-squashing measure OSQ.

B Summary of spectral properties on graphs

Basic notions of spectral theory on graphs.

Throughout the appendix, we let Δ be the normalized graph Laplacian defined by $\Delta = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$. It is known [17] that the graph Laplacian is a symmetrically, positive semi-definite matrix whose spectral decomposition takes the form

$$\Delta = \sum_{\ell=0}^{n-1} \lambda_{\ell} \phi_{\ell} \phi_{\ell}^{\top}, \quad (9)$$

where $\{\phi_{\ell}\}$ is an orthonormal basis in \mathbb{R}^n and $0 = \lambda_0 < \lambda_1 < \dots < \lambda_{n-1}$ – recall that since we assume G to be connected, the zero eigenvalue has multiplicity one, i.e. $\lambda_1 > 0$. We also note that we typically write $\phi_{\ell}(v)$ for the value of ϕ_{ℓ} at $v \in V$, and that the kernel of Δ is spanned by ϕ_0 with $\phi_0(v) = \sqrt{d_v/2|E|}$. The results would extend to the bipartite case as usual when doing spectral analysis one too if graphs, we exclude the edge

case of the bipartite graph to make sure that the largest eigenvalue of the graph Laplacian satisfies $\lambda_{n-1} < 2$ – yet all results hold for the bipartite case too provided we take $\|\nabla_2\psi\| < 1$. Finally, we let Δ^\dagger denote the pseudo-inverse of the graph Laplacian, which can be written as

$$\Delta^\dagger = \sum_{\ell=1}^{n-1} \frac{1}{\lambda_\ell} \phi_\ell \phi_\ell^\top, \quad (10)$$

and we emphasize that the sum starts from $\ell = 1$ since we need to ignore the kernel of Δ spanned by the orthonormal vector ϕ_0 .

Basic properties of Random Walks on graphs. A simple Random Walk (RW) on G is a Markov chain supported on the nodes V with transition matrix defined by $\mathbf{P}(v, u) = d_v^{-1}$. While a RW can be studied through different properties, the one we are interested in is the *commute time* τ , which represents the expected number of steps for a RW starting at v , to visit u and then come back to v . The commute time is a *distance* on the graph and captures the diffusion properties associated with the underlying topology. In fact, while nodes that are distant often have larger commute time, the latter is more expressive than the shortest-walk graph-distance, since it also accounts (for example) for the number of paths connecting two given nodes. Thanks to [38], we can write down the commute time among two nodes using the spectral representation of the graph Laplacian in (9):

$$\tau(v, u) = 2|E| \sum_{\ell=1}^{n-1} \frac{1}{\lambda_\ell} \left(\frac{\phi_\ell(v)}{\sqrt{d_v}} - \frac{\phi_\ell(u)}{\sqrt{d_u}} \right)^2. \quad (11)$$

C Proofs and additional details of Section 3

The goal of this section amounts to proving Theorem 3.2. To work towards this result, we first derive bounds on the Jacobian and Hessian of a single node feature after m layers before the readout READ operation. We emphasize that our analysis below is novel, when compared to previous works of [8, 24], on many accounts. First, [8, 24] do not consider higher (second) order derivatives, limiting their discussion to the case of first order derivatives, which are not suited to capture notions of mixing among features – we will expand on this topic in Section E. Second, even for the case of first-order derivatives, our result below is more general since it holds for all MPNNs as in (2), which includes (i) message-functions ψ that also depend on the input features (as for GatedGCN), and (ii) choices of message-passing matrices \mathbf{A} that could be weighted and (or) asymmetric. Third, the analysis in [8, 24] does not account for the role of the readout map and hence fails to study the expressive power of graph-level prediction of MPNNs as measured by the mixing they generate among nodes.

Conventions and notations for the proofs. First, we recall that $\mathbf{h}_v^{(0)} = \mathbf{x}_v \in \mathbb{R}^d$ is the input feature at node v . Below, we write $h_v^{(t),\alpha}$ for the α -th entry of the feature $\mathbf{h}_v^{(t)}$. To simplify the notations, we rewrite the layer-update in (2) using coordinates as

$$h_v^{(t),\alpha} = \sigma(\tilde{h}_v^{(t-1),\alpha}), \quad 1 \leq \alpha \leq d, \quad (12)$$

where $\tilde{h}_v^{(t-1),\alpha}$ is the entry α of the pre-activated feature of node v at layer t . We also let $\partial_{1,p}\psi^{(t),r}$ and $\partial_{2,p}\psi^{(t),r}$ be the p -th derivative of $(\psi^{(t)}(\cdot, x))_r$ and of $(\psi^{(t)}(x, \cdot))_r$, respectively. To avoid cumbersome notations, we usually omit to write the arguments of the derivatives of the message-functions ψ . Similarly, we let $\nabla_1\psi$ ($\nabla_2\psi$) be the $d \times d$ Jacobian matrix of ψ with respect to the first (second) variable. Finally, given nodes $i, v, u \in V$ we introduce the following terms:

$$\nabla_u \mathbf{h}_v^{(m)} := \frac{\partial \mathbf{h}_v^{(m)}}{\partial \mathbf{x}_u} \in \mathbb{R}^{d \times d}, \quad \nabla_{uv}^2 \mathbf{h}_i^{(m)} := \frac{\partial^2 \mathbf{h}_i^{(m)}}{\partial \mathbf{x}_u \partial \mathbf{x}_v} \in \mathbb{R}^{d \times (d \times d)}.$$

First, we derive an upper bound on the first-order derivatives of the node-features. This will provide useful to derive the more general second-order estimate of the MPNN-prediction. We highlight that the result below

extends the analysis in [24] to MPNNs with arbitrary (i.e. non-linear) message functions ψ , such as GatedGCN [10].

Theorem C.1. *Given MPNNs as in (2), let σ and $\psi^{(t)}$ be \mathcal{C}^1 functions and assume $|\sigma'| \leq c_\sigma$, $\|\boldsymbol{\Omega}^{(t)}\| \leq \omega$, $\|\mathbf{W}^{(t)}\| \leq \mathbf{w}$, $\|\nabla_1 \psi^{(t)}\| \leq c_1$, and $\|\nabla_2 \psi^{(t)}\| \leq c_2$. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be*

$$\mathbf{S} := \frac{\omega}{\mathbf{w}} \mathbf{I} + c_1 \text{diag}(\mathbf{A}\mathbf{1}) + c_2 \mathbf{A}.$$

Given nodes $v, u \in \mathcal{V}$ and m the number of layers, the following holds:

$$\|\nabla_u \mathbf{h}_v^{(m)}\| \leq (c_\sigma \mathbf{w})^m (\mathbf{S}^m)_{vu}. \quad (13)$$

Proof. Recall that the dimension of the features is taken to be d for any layer $1 \leq t \leq m$. We proceed by induction. If $m = 1$ and we fix entries $1 \leq \alpha, \beta \leq d$, then using the shorthand in (12), we obtain

$$\begin{aligned} (\nabla_u \mathbf{h}_v^{(1)})_{\alpha\beta} &= \sigma'(\tilde{h}_v^{(0),\alpha}) \left(\Omega_{\alpha\beta}^{(1)} \delta_{vu} + \sum_r W_{\alpha r}^{(1)} \sum_j A_{vj} \left(\partial_{1,\beta} \psi^{(1),r} \delta_{vu} + \partial_{2,\beta} \psi^{(1),r} \delta_{ju} \right) \right) \\ &= \left(\text{diag}(\sigma'(\tilde{\mathbf{h}}_v^{(0)})) \left(\boldsymbol{\Omega}^{(1)} \delta_{vu} + \mathbf{W}^{(1)} \left(\sum_j A_{vj} \delta_{vu} \nabla_1 \psi^{(1)} + A_{vu} \nabla_2 \psi^{(1)} \right) \right) \right)_{\alpha\beta}. \end{aligned}$$

Therefore, we can bound the (spectral) norm of the Jacobian on the left hand side by

$$\begin{aligned} \|\nabla_u \mathbf{h}_v^{(1)}\| &\leq \|\text{diag}(\sigma'(\tilde{\mathbf{h}}_v^{(0)}))\| \left(\|\boldsymbol{\Omega}^{(1)}\| \delta_{vu} + \|\mathbf{W}^{(1)}\| (c_1 \sum_j A_{vj} \delta_{vu} + c_2 A_{vu}) \right) \\ &\leq c_\sigma (\omega \delta_{vu} + \mathbf{w} (c_1 \sum_j A_{vj} \delta_{vu} + c_2 A_{vu})) = c_\sigma \mathbf{w} \mathbf{S}_{vu}, \end{aligned}$$

which proves the estimate on the Jacobian for the case of $m = 1$. We now take the induction step, and follow the same argument above to write the node Jacobian after m layers as

$$\begin{aligned} (\nabla_u \mathbf{h}_v^{(m)})_{\alpha\beta} &= \sigma'(\tilde{h}_v^{(m-1),\alpha}) \left(\sum_r \Omega_{\alpha r}^{(m)} (\nabla_u \mathbf{h}_v^{(m-1)})_{r\beta} \right) \\ &\quad + \sigma'(\tilde{h}_v^{(m-1),\alpha}) \left(W_{\alpha r}^{(m)} \sum_j A_{vj} \sum_p \left(\partial_{1,p} \psi^{(m),r} (\nabla_u \mathbf{h}_v^{(m-1)})_{p\beta} + \partial_{2,p} \psi^{(m),r} (\nabla_u \mathbf{h}_j^{(m-1)})_{p\beta} \right) \right) \\ &= \left(\text{diag}(\sigma'(\tilde{\mathbf{h}}_v^{(m-1)})) \boldsymbol{\Omega}^{(m)} \nabla_u \mathbf{h}_v^{(m-1)} \right)_{\alpha\beta} \\ &\quad + \left(\text{diag}(\sigma'(\tilde{\mathbf{h}}_v^{(m-1)})) \mathbf{W}^{(m)} \left(\sum_j A_{vj} \nabla_1 \psi^{(m)} \nabla_u \mathbf{h}_v^{(m-1)} + A_{vj} \nabla_2 \psi^{(m)} \nabla_u \mathbf{h}_j^{(m-1)} \right) \right)_{\alpha\beta}. \end{aligned}$$

Therefore, we can use the induction step to bound the Jacobian as

$$\begin{aligned} \|\nabla_u \mathbf{h}_v^{(m)}\| &\leq c_\sigma \omega (c_\sigma \mathbf{w})^{m-1} (\mathbf{S}^{m-1})_{vu} + (c_\sigma \mathbf{w})^m \left(\sum_j A_{vj} c_1 (\mathbf{S}^{m-1})_{vu} + A_{vj} c_2 (\mathbf{S}^{m-1})_{ju} \right) \\ &= (c_\sigma \mathbf{w})^m \left(\left(\frac{\omega}{\mathbf{w}} \mathbf{I} + c_1 \text{diag}(\mathbf{A}\mathbf{1}) + c_2 \mathbf{A} \right) (\mathbf{S}^{m-1}) \right)_{vu} = (c_\sigma \mathbf{w})^m (\mathbf{S}^m)_{vu}, \end{aligned}$$

which completes the proof for the first-order bounds. \square

Before we move to the second-order estimates, we introduce some additional preliminary notations. Given nodes i, v, u , a matrix $\mathbf{S} \in \mathbb{R}^{n \times n}$ – which will always be chosen as per (5) – and an integer ℓ , we write

$$\mathbf{P}_{i(vu)}^{(\ell)} := (\mathbf{S}^\ell)_{iv} (\mathbf{A}\mathbf{S}^\ell)_{iu} + (\mathbf{S}^\ell)_{iu} (\mathbf{A}\mathbf{S}^\ell)_{iv} + \sum_j (\mathbf{S}^\ell)_{jv} (\text{diag}(\mathbf{A}\mathbf{1}) + \mathbf{A})_{ij} (\mathbf{S}^\ell)_{ju}. \quad (14)$$

In particular, we denote by $\mathbf{P}_{(vu)}^{(\ell)} \in \mathbb{R}^n$ the vector with entries $(\mathbf{P}_{(vu)}^{(\ell)})_i = \mathbf{P}_{i(vu)}^{(\ell)}$, for $1 \leq i \leq n$.

Theorem C.2. Given MPNNs as in (2), let σ and $\psi^{(t)}$ be \mathcal{C}^2 functions and assume $|\sigma'|, |\sigma''| \leq c_\sigma$, $\|\boldsymbol{\Omega}^{(t)}\| \leq \omega$, $\|\mathbf{W}^{(t)}\| \leq \mathbf{w}$, $\|\nabla_1 \psi^{(t)}\| \leq c_1$, $\|\nabla_2 \psi^{(t)}\| \leq c_2$, $\|\nabla^2 \psi^{(t)}\| \leq c^{(2)}$. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be

$$\mathbf{S} := \frac{\omega}{\mathbf{w}} \mathbf{I} + c_1 \text{diag}(\mathbf{A1}) + c_2 \mathbf{A}.$$

Given nodes $i, v, u \in \mathbf{V}$, if $\mathbf{P}_{(vu)}^{(\ell)} \in \mathbb{R}^n$ is as in (14) and m is the number of layers, then we derive

$$\begin{aligned} \|\nabla_{uv}^2 \mathbf{h}_i^{(m)}\| &\leq \sum_{k=0}^{m-1} \sum_{j \in \mathbf{V}} (c_\sigma \mathbf{w})^{2m-k-1} \mathbf{w} (\mathbf{S}^{m-k})_{jv} (\mathbf{S}^k)_{ij} (\mathbf{S}^{m-k})_{ju} \\ &+ c^{(2)} \sum_{\ell=0}^{m-1} (c_\sigma \mathbf{w})^{m+\ell} (\mathbf{S}^{m-1-\ell} \mathbf{P}_{(vu)}^{(\ell)})_i. \end{aligned} \quad (15)$$

Proof. First, we note that $\nabla_{uv}^2 \mathbf{h}_i^{(m)}$ is a matrix of dimension $\mathbb{R}^{d \times (d \times d)}$. We then use the following ordering for indexing the columns – which is consistent with a typical way of labelling columns of the Kronecker product of matrices, as detailed below (note that indices here start from 1):

$$\frac{\partial^2 h_i^{(m), \alpha}}{\partial x_v^\beta \partial x_u^\gamma} := \left(\nabla_{uv}^2 \mathbf{h}_i^{(m)} \right)_{\alpha, d(\beta-1)+\gamma}. \quad (16)$$

As above, we proceed by induction and start from the case $m = 1$:

$$\begin{aligned} \left(\nabla_{uv}^2 \mathbf{h}_i^{(1)} \right)_{\alpha, d(\beta-1)+\gamma} &= \sigma''(\tilde{h}_i^{(0), \alpha}) \left(\Omega_{\alpha\gamma}^{(1)} \delta_{iv} + \sum_r W_{\alpha r}^{(1)} \sum_j \mathbf{A}_{ij} (\delta_{iv} \partial_{1,\gamma} \psi^{(1),r} + \delta_{jv} \partial_{2,\gamma} \psi^{(1),r}) \right) \\ &\times \left(\Omega_{\alpha\beta}^{(1)} \delta_{iu} + \sum_r W_{\alpha r}^{(1)} \sum_j \mathbf{A}_{ij} (\delta_{iu} \partial_{1,\beta} \psi^{(1),r} + \delta_{ju} \partial_{2,\beta} \psi^{(1),r}) \right) \\ &+ \sigma'(\tilde{h}_i^{(0), \alpha}) \sum_r W_{\alpha r}^{(1)} \left(\sum_j \mathbf{A}_{ij} \delta_{iu} (\partial_{1,\gamma} \partial_{1,\beta} \psi^{(1),r} \delta_{iv} + \partial_{2,\gamma} \partial_{1,\beta} \psi^{(1),r} \delta_{jv}) \right) \\ &+ \sigma'(\tilde{h}_i^{(0), \alpha}) \sum_r W_{\alpha r}^{(1)} \left(\mathbf{A}_{iu} (\partial_{1,\gamma} \partial_{2,\beta} \psi^{(1),r} \delta_{iv} + \partial_{2,\gamma} \partial_{2,\beta} \psi^{(1),r} \delta_{uv}) \right) \\ &:= (Q_1)_{\alpha,\beta,\gamma} + (Q_2)_{\alpha,\beta,\gamma} + (Q_3)_{\alpha,\beta,\gamma}, \end{aligned}$$

where Q_1 is the term containing second derivatives of ψ while Q_2, Q_3 are the remaining expressions including second-order derivatives of the message functions ψ . Using the same strategy as for the first-order estimates, we can rewrite the first term Q_1 as

$$\begin{aligned} (Q_1)_{\alpha,\beta,\gamma} &= \left(\text{diag}(\sigma''(\tilde{\mathbf{h}}_v^{(0)})) \left(\boldsymbol{\Omega}^{(1)} \delta_{iv} + \mathbf{W}^{(1)} \left(\sum_j \mathbf{A}_{ij} \delta_{iv} \nabla_1 \psi^{(1)} + \mathbf{A}_{iv} \nabla_2 \psi^{(1)} \right) \right) \right)_{\alpha\gamma} \\ &\times \left(\boldsymbol{\Omega}^{(1)} \delta_{iu} + \mathbf{W}^{(1)} \left(\sum_j \mathbf{A}_{ij} \delta_{iu} \nabla_1 \psi^{(1)} + \mathbf{A}_{iu} \nabla_2 \psi^{(1)} \right) \right)_{\alpha\beta} \end{aligned}$$

We now observe that given two matrices $\mathbf{B}, \mathbf{C} \in \mathbb{R}^{d \times d}$ and $1 \leq \alpha, \alpha', \beta, \gamma \leq d$, the entries of the Kronecker product $\mathbf{B} \otimes \mathbf{C}$ can be indexed as

$$(\mathbf{B} \otimes \mathbf{C})_{d(\alpha-1)+\alpha', d(\beta-1)+\gamma} = B_{\alpha\beta} C_{\alpha'\gamma}.$$

We now introduce the $d \times (d \times d)$ sub-matrix of $\mathbf{B} \otimes \mathbf{C}$ defined by

$$(\mathbf{B} \otimes \mathbf{C})'_{\alpha, d(\beta-1)+\gamma} = B_{\alpha\beta} C_{\alpha\gamma}. \quad (17)$$

Therefore, we can rewrite $(Q_1)_{\alpha,\beta,\gamma}$ as the entry $(\alpha, d(\beta-1)+\gamma)$ of the $d \times (d \times d)$ sub-matrix

$$(\mathbf{Q}_1)_{\alpha, d(\beta-1)+\gamma} = (\mathbf{B} \otimes \mathbf{C})'_{\alpha, d(\beta-1)+\gamma}, \quad (18)$$

where

$$\begin{aligned}\mathbf{B} &:= \text{diag}(\sigma''(\tilde{\mathbf{h}}_v^{(0)})) \left(\boldsymbol{\Omega}^{(1)} \delta_{iv} + \mathbf{W}^{(1)} \left(\sum_j \mathbf{A}_{ij} \delta_{iv} \nabla_1 \psi^{(1)} + \mathbf{A}_{iv} \nabla_2 \psi^{(1)} \right) \right), \\ \mathbf{C} &:= \boldsymbol{\Omega}^{(1)} \delta_{iu} + \mathbf{W}^{(1)} \left(\sum_j \mathbf{A}_{ij} \delta_{iu} \nabla_1 \psi^{(1)} + \mathbf{A}_{iu} \nabla_2 \psi^{(1)} \right).\end{aligned}$$

Next, we proceed to write $(Q_2)_{\alpha, \beta, \gamma}$ in matricial form. Before we do that, we observe that the Hessian of the message functions $(\mathbf{x}_i, \mathbf{x}_j) \mapsto \psi^{(t)}(\mathbf{x}_i, \mathbf{x}_j)$ takes the form

$$\nabla^2 \psi^{(t)} = \begin{pmatrix} \nabla_{11}^2 \psi^{(t)} & \nabla_{12}^2 \psi^{(t)} \\ \nabla_{21}^2 \psi^{(t)} & \nabla_{22}^2 \psi^{(t)} \end{pmatrix},$$

where $\nabla_{ab}^2 \psi^{(t)} \in \mathbb{R}^{d \times (d \times d)}$ and is indexed as follows

$$(\nabla_{ab}^2 \psi^{(t)})_{r, d(\beta-1)+\gamma} = \partial_{a,\beta} \partial_{b,\gamma} \psi^{(t),r},$$

where $a, b \in \{1, 2\}$. Using these notations, we note that

$$\sum_r W_{\alpha r}^{(1)} \partial_{1,\gamma} \partial_{1,\beta} \psi^{(1),r} = \left(\mathbf{W}^{(1)} \nabla_{11}^2 \psi^{(1)} \right)_{\alpha, d(\beta-1)+\gamma}.$$

Therefore, we derive

$$\begin{aligned}(Q_2)_{\alpha, \beta, \gamma} &= (\mathbf{Q}_2)_{\alpha, d(\beta-1)+\gamma} = \sum_j \mathbf{A}_{ij} \delta_{iu} \delta_{iv} (\text{diag}(\sigma'(\tilde{\mathbf{h}}_i^{(0)})) \mathbf{W}^{(1)} \nabla_{11}^2 \psi^{(1)})_{\alpha, d(\beta-1)+\gamma} \\ &\quad + \mathbf{A}_{iv} \delta_{iu} (\text{diag}(\sigma'(\tilde{\mathbf{h}}_i^{(0)})) \mathbf{W}^{(1)} \nabla_{12}^2 \psi^{(1)})_{\alpha, d(\beta-1)+\gamma}.\end{aligned}\tag{19}$$

A similar argument works for Q_3 :

$$\begin{aligned}(Q_3)_{\alpha, \beta, \gamma} &= (\mathbf{Q}_3)_{\alpha, d(\beta-1)+\gamma} = \mathbf{A}_{iu} \delta_{iv} (\text{diag}(\sigma'(\tilde{\mathbf{h}}_i^{(0)})) \mathbf{W}^{(1)} \nabla_{21}^2 \psi^{(1)})_{\alpha, d(\beta-1)+\gamma} \\ &\quad + \mathbf{A}_{iu} \delta_{uv} (\text{diag}(\sigma'(\tilde{\mathbf{h}}_i^{(0)})) \mathbf{W}^{(1)} \nabla_{22}^2 \psi^{(1)})_{\alpha, d(\beta-1)+\gamma}.\end{aligned}\tag{20}$$

Therefore, we can combine (18), (19), and (20) to write

$$\begin{aligned}\|\nabla_{uv}^2 \mathbf{h}_i^{(1)}\| &\leq \|\mathbf{Q}_1\| + \|\mathbf{Q}_2\| + \|\mathbf{Q}_3\| \\ &\leq c_\sigma (\omega \delta_{iv} + w(c_1 \text{diag}(\mathbf{A}_1)_i \delta_{iv} + c_2 \mathbf{A}_{iv})) (\omega \delta_{iu} + w(c_1 \text{diag}(\mathbf{A}_1)_i \delta_{iu} + c_2 \mathbf{A}_{iu})) \\ &\quad + c_\sigma w c^{(2)} (\text{diag}(\mathbf{A}_1)_i \delta_{iv} \delta_{iu} + \mathbf{A}_{iv} \delta_{iu}) \\ &\quad + c_\sigma w c^{(2)} (\mathbf{A}_{iu} \delta_{iv} + \mathbf{A}_{iu} \delta_{uv}).\end{aligned}$$

Finally, we can rely on (14) to re-arrange the equation above as

$$\begin{aligned}\|\nabla_{uv}^2 \mathbf{h}_i^{(1)}\| &\leq (c_\sigma w) (w(\mathbf{S})_{iv} (\mathbf{S})_{iu}) + w c^{(2)} c_\sigma (\delta_{iv} \mathbf{A}_{iu} + \delta_{iu} \mathbf{A}_{iv} + \sum_j \delta_{jv} (\text{diag}(\mathbf{A}_1) + \mathbf{A})_{ij} \delta_{ju}) \\ &= (c_\sigma w) (w(\mathbf{S})_{iv} (\mathbf{S})_{iu}) + c^{(2)} c_\sigma w \mathbf{P}_{i(vu)}^{(0)},\end{aligned}$$

which proves the bound for the second-order derivatives in the case $m = 1$.

We now assume that the claim holds for all layers $t \leq m - 1$, and compute the second order derivative after m

layers:

$$\begin{aligned}
\left(\nabla_{uv}^2 \mathbf{h}_i^{(m)}\right)_{\alpha, d(\beta-1)+\gamma} &= \sigma''(\tilde{h}_i^{(m-1), \alpha}) \\
&\times \left(\sum_r \Omega_{\alpha r}^{(m)} (\nabla_u \mathbf{h}_i^{(m-1)})_{r\beta}\right) \\
&+ W_{\alpha r}^{(m)} \sum_j A_{ij} \left(\sum_p \partial_{1,p} \psi^{(m), r} (\nabla_u \mathbf{h}_i^{(m-1)})_{p\beta} + \partial_{2,p} \psi^{(m), r} (\nabla_u \mathbf{h}_j^{(m-1)})_{p\beta}\right) \\
&\times \left(\sum_r \Omega_{\alpha r}^{(m)} (\nabla_v \mathbf{h}_i^{(m-1)})_{r\gamma}\right) \\
&+ W_{\alpha r}^{(m)} \sum_j A_{ij} \left(\sum_p \partial_{1,p} \psi^{(m), r} (\nabla_v \mathbf{h}_i^{(m-1)})_{p\gamma} + \partial_{2,p} \psi^{(m), r} (\nabla_v \mathbf{h}_j^{(m-1)})_{p\gamma}\right) \\
&+ \sigma'(\tilde{h}_i^{(m-1), \alpha}) \sum_r W_{\alpha r}^{(m)} \sum_j A_{ij} \sum_{p,q} \partial_{1,p} \partial_{1,q} \psi^{(m), r} (\nabla_u \mathbf{h}_i^{(m-1)})_{p\beta} (\nabla_v \mathbf{h}_i^{(m-1)})_{q\gamma} \\
&+ \sigma'(\tilde{h}_i^{(m-1), \alpha}) \sum_r W_{\alpha r}^{(m)} \sum_j A_{ij} \sum_{p,q} \partial_{1,p} \partial_{2,q} \psi^{(m), r} (\nabla_u \mathbf{h}_i^{(m-1)})_{p\beta} (\nabla_v \mathbf{h}_j^{(m-1)})_{q\gamma} \\
&+ \sigma'(\tilde{h}_i^{(m-1), \alpha}) \sum_r W_{\alpha r}^{(m)} \sum_j A_{ij} \sum_{p,q} \partial_{1,p} \partial_{2,q} \psi^{(m), r} (\nabla_u \mathbf{h}_j^{(m-1)})_{q\beta} (\nabla_v \mathbf{h}_i^{(m-1)})_{p\gamma} \\
&+ \sigma'(\tilde{h}_i^{(m-1), \alpha}) \sum_r W_{\alpha r}^{(m)} \sum_j A_{ij} \sum_{p,q} \partial_{2,p} \partial_{2,q} \psi^{(m), r} (\nabla_u \mathbf{h}_j^{(m-1)})_{p\beta} (\nabla_v \mathbf{h}_j^{(m-1)})_{q\gamma} \\
&+ \sigma'(\tilde{h}_i^{(m-1), \alpha}) \sum_r \Omega_{\alpha r}^{(m)} (\nabla_{uv}^2 \mathbf{h}_i^{(m-1)})_{r, d(\beta-1)+\gamma} \\
&+ \sigma'(\tilde{h}_i^{(m-1), \alpha}) \sum_r W_{\alpha r}^{(m)} \sum_j A_{ij} \left(\sum_p \partial_{1,p} \psi^{(m), r} (\nabla_{uv}^2 \mathbf{h}_i^{(m-1)})_{p, d(\beta-1)+\gamma}\right) \\
&+ \sigma'(\tilde{h}_i^{(m-1), \alpha}) \sum_r W_{\alpha r}^{(m)} \sum_j A_{ij} \left(\sum_p \partial_{2,p} \psi^{(m), r} (\nabla_{uv}^2 \mathbf{h}_j^{(m-1)})_{p, d(\beta-1)+\gamma}\right) \\
&:= R_{\alpha, \beta, \gamma} + \sum_{a, b \in \{1, 2\}} (Q_{ab})_{\alpha, \beta, \gamma} + Z_{\alpha, \beta, \gamma},
\end{aligned}$$

where R is the term containing second derivatives of the non-linear map σ , Q_{ab} is indexed according to the second derivatives of the message-functions ψ , and finally Z is the term containing second-order derivatives of the features. For the term $R_{\alpha, \beta, \gamma}$ we can argue as in the $m = 1$ case and use the sub-matrix notation in (17) to rewrite it as the entry $(\alpha, d(\beta - 1) + \gamma)$ of the $d \times (d \times d)$ sub-matrix

$$\mathbf{R}_{\alpha, d(\beta-1)+\gamma} = (\mathbf{B} \otimes \mathbf{C})'_{\alpha, d(\beta-1)+\gamma}, \quad (21)$$

where

$$\begin{aligned}
\mathbf{B} &:= \text{diag}(\sigma''(\tilde{\mathbf{h}}_i^{(m-1)})) (\boldsymbol{\Omega}^{(m)} \nabla_u \mathbf{h}_i^{(m-1)} + \mathbf{W}^{(m)} (\sum_j A_{ij} \nabla_1 \psi^{(m)} \nabla_u \mathbf{h}_i^{(m-1)} + \nabla_2 \psi^{(m)} \nabla_u \mathbf{h}_j^{(m-1)})), \\
\mathbf{C} &:= \boldsymbol{\Omega}^{(m)} \nabla_v \mathbf{h}_i^{(m-1)} + \mathbf{W}^{(m)} (\sum_j A_{ij} \nabla_1 \psi^{(m)} \nabla_v \mathbf{h}_i^{(m-1)} + \nabla_2 \psi^{(m)} \nabla_v \mathbf{h}_j^{(m-1)})
\end{aligned}$$

Next we consider the terms $(Q_{ab})_{\alpha, \beta, \gamma}$. Without loss of generality, we focus on $(Q_{11})_{\alpha, \beta, \gamma}$ and use again the same argument in the $m = 1$ case, to rewrite it as $(Q_{11})_{\alpha, \beta, \gamma} = (\mathbf{Q}_{11})_{\alpha, d(\beta-1)+\gamma}$ where

$$\mathbf{Q}_{11} = \text{diag}(\sigma'(\tilde{\mathbf{h}}_i^{(m-1)})) \sum_j A_{ij} (\mathbf{W}^{(m)} \nabla_{11}^2 \psi^{(m)} \nabla_u \mathbf{h}_i^{(m-1)} \otimes \nabla_v \mathbf{h}_i^{(m-1)}), \quad (22)$$

where again we are indexing the matrix $\nabla_{11}^2 \psi^{(m)}$ by

$$(\nabla_{11}^2 \psi^{(m)})_{r, p(d-1)+q} = \partial_{1,p} \partial_{1,q} \psi^{(m), r}.$$

The other Q -terms can be estimated similarly. Finally, we rewrite $Z_{\alpha,\beta,\gamma} = (\mathbf{Z})_{\alpha,d(\beta-1)+\gamma}$, where

$$\mathbf{Z} = \text{diag}(\sigma'(\tilde{\mathbf{h}}_i^{(m-1)})) \left(\boldsymbol{\Omega}^{(m)} \nabla_{uv}^2 \mathbf{h}_i^{(m-1)} + \mathbf{W}^{(m)} \sum_j \mathbf{A}_{ij} (\nabla_1 \psi^{(m)} \nabla_{uv}^2 \mathbf{h}_i^{(m-1)} + \nabla_2 \psi^{(m)} \nabla_{uv}^2 \mathbf{h}_j^{(m-1)}) \right) \quad (23)$$

Therefore, we have rewritten the second-derivatives of the features in matricial form as

$$\nabla_{uv}^2 \mathbf{h}_i^{(m)} = \mathbf{R} + \sum_{a,b \in \{1,2\}} \mathbf{Q}_{ab} + \mathbf{Z}.$$

To complete the proof, we now simply need to estimate the three terms and show they fit the recursion claimed for m . For the case of \mathbf{R} in (21), we find

$$\begin{aligned} \|\mathbf{R}\| &\leq c_\sigma (\omega \|\nabla_u \mathbf{h}_i^{(m-1)}\| + \mathbf{w} (c_1 \text{diag}(\mathbf{A}\mathbf{1})_i \|\nabla_u \mathbf{h}_i^{(m-1)}\| + c_2 \sum_j \mathbf{A}_{ij} \|\nabla_u \mathbf{h}_j^{(m-1)}\|)) \\ &\quad \times (\omega \|\nabla_v \mathbf{h}_i^{(m-1)}\| + \mathbf{w} (c_1 \text{diag}(\mathbf{A}\mathbf{1})_i \|\nabla_v \mathbf{h}_i^{(m-1)}\| + c_2 \sum_j \mathbf{A}_{ij} \|\nabla_v \mathbf{h}_j^{(m-1)}\|)). \end{aligned}$$

If we write $D\mathbf{h}^{(m-1)} \in \mathbb{R}^{n \times n}$ as the matrix with entries $(D\mathbf{h}^{(m-1)})_{ij} = \|\nabla_j \mathbf{h}_i^{(m-1)}\|$, then we obtain

$$\|\mathbf{R}\| \leq c_\sigma \mathbf{w} (\mathbf{w} \mathbf{S} D\mathbf{h}^{(m-1)})_{iv} (\mathbf{S} D\mathbf{h}^{(m-1)})_{iu}.$$

We can then plug the first-order estimates derived in Theorem C.1 and obtain

$$\|\mathbf{R}\| \leq c_\sigma \mathbf{w} (\mathbf{w} \mathbf{S} (c_\sigma \mathbf{w})^{m-1} \mathbf{S}^{m-1})_{iv} (\mathbf{S} (c_\sigma \mathbf{w})^{m-1} \mathbf{S}^{m-1})_{iu} = (c_\sigma \mathbf{w})^{2m-1} (\mathbf{w} (\mathbf{S}^m)_{iv} (\mathbf{S}^m)_{iu}). \quad (24)$$

Next, we move onto the Q -terms, and use again the first-order estimates in Theorem C.1 – and the fact that we can bound the norm of $\nabla_{ab}^2 \psi^{(m)}$ by $c^{(2)}$ – to derive

$$\begin{aligned} \left\| \sum_{a,b \in \{1,2\}} \mathbf{Q}_{ab} \right\| &\leq c^{(2)} (c_\sigma \mathbf{w})^{2m-1} (\text{diag}(\mathbf{A}\mathbf{1})_i (\mathbf{S}^{m-1})_{iv} (\mathbf{S}^{m-1})_{iu} + \sum_j \mathbf{A}_{ij} (\mathbf{S}^{m-1})_{ju} (\mathbf{S}^{m-1})_{jv}) \\ &\quad + c^{(2)} (c_\sigma \mathbf{w})^{2m-1} ((\mathbf{S}^{m-1})_{iv} (\mathbf{A} \mathbf{S}^{m-1})_{iu} + (\mathbf{S}^{m-1})_{iu} (\mathbf{A} \mathbf{S}^{m-1})_{iv}) \\ &= c^{(2)} (c_\sigma \mathbf{w})^{2m-1} \mathbf{P}_{i(vu)}^{(m-1)}. \end{aligned} \quad (25)$$

Finally, if we let $D^2 \mathbf{h}_{vu} \in \mathbb{R}^n$ be the vector with entries $(D^2 \mathbf{h}_{vu})_i = \|\nabla_{uv}^2 \mathbf{h}_i^{(m-1)}\|$, then

$$\begin{aligned} \|\mathbf{Z}\| &\leq c_\sigma \left(\omega \|\nabla_{uv}^2 \mathbf{h}_i^{(m-1)}\| + \mathbf{w} \left(c_1 \text{diag}(\mathbf{A}\mathbf{1})_i \|\nabla_{uv}^2 \mathbf{h}_i^{(m-1)}\| + c_2 \sum_j \mathbf{A}_{ij} \|\nabla_{uv}^2 \mathbf{h}_j^{(m-1)}\| \right) \right) \\ &= c_\sigma \mathbf{w} (\mathbf{S} D^2 \mathbf{h}_{vu})_i. \end{aligned} \quad (26)$$

Therefore, we can use the induction to derive

$$\begin{aligned}
\|\mathbf{Z}\| &\leq c_\sigma \mathbf{w} \sum_s \mathbf{S}_{is} \sum_{k=0}^{m-2} \sum_{j \in \mathbf{V}} (c_\sigma \mathbf{w})^{2m-2-k-1} \mathbf{w}(\mathbf{S}^{m-1-k})_{jv} (\mathbf{S}^k)_{sj} (\mathbf{S}^{m-1-k})_{ju} \\
&+ c_\sigma \mathbf{w} \sum_s \mathbf{S}_{is} (c^{(2)} \sum_{\ell=0}^{m-2} (c_\sigma \mathbf{w})^{m-1+\ell} (\mathbf{S}^{m-2-\ell} \mathbf{P}_{(vu)}^{(\ell)})_s) \\
&= \sum_{k=0}^{m-2} \sum_{j \in \mathbf{V}} (c_\sigma \mathbf{w})^{2m-k-2} \mathbf{w}(\mathbf{S}^{m-1-k})_{jv} (\mathbf{S}^{k+1})_{ij} (\mathbf{S}^{m-1-k})_{ju} \\
&+ c^{(2)} \left(\sum_{\ell=0}^{m-2} (c_\sigma \mathbf{w})^{m+\ell} (\mathbf{S}^{m-1-\ell} \mathbf{P}_{(vu)}^{(\ell)})_i \right) \\
&= \sum_{k=1}^{m-1} \sum_{j \in \mathbf{V}} (c_\sigma \mathbf{w})^{2m-k-1} \mathbf{w}(\mathbf{S}^{m-k})_{jv} (\mathbf{S}^k)_{ij} (\mathbf{S}^{m-k})_{ju} + c^{(2)} \left(\sum_{\ell=0}^{m-2} (c_\sigma \mathbf{w})^{m+\ell} (\mathbf{S}^{m-1-\ell} \mathbf{P}_{(vu)}^{(\ell)})_i \right)
\end{aligned}$$

By (24), we derive that the \mathbf{R} -term corresponds to the $k = 0$ entry of the first sum, while (25) corresponds to the case $\ell = m - 1$ of the second sum, which completes the induction and hence our proof. \square

C.1 Proof of Theorem 3.2

We can now use the previous characterization to derive estimates on the Hessian of the graph-level function computed by MPNNs. We restate Theorem 3.2 here for convenience.

Theorem 3.2. *Consider an MPNN of depth m as in (2), where σ and $\psi^{(t)}$ are \mathcal{C}^2 functions and we denote the bounds on their derivatives and on the norm of the weights as above. Let \mathbf{S} and \mathbf{Q}_k be defined as in (5) and (6), respectively. If the readout is MAX, MEAN or SUM and $\boldsymbol{\theta}$ in (3) has unit norm, then the mixing $\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u)$ induced by the MPNN over the features of nodes v, u satisfies*

$$\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u) \leq \sum_{k=0}^{m-1} (c_\sigma \mathbf{w})^{2m-k-1} \left(\mathbf{w}(\mathbf{S}^{m-k})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k \mathbf{S}^{m-k} + c^{(2)} \mathbf{Q}_k)_{vu} \right). \quad (7)$$

Proof. First, we recall that according to Definition 3.1, we are interested in bounding the quantity

$$\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u) = \max_{\mathbf{X}} \max_{1 \leq \beta, \gamma \leq d} \left| \frac{\partial^2 y_{\mathbf{G}}^{(m)}(\mathbf{X})}{\partial x_u^\beta \partial x_v^\gamma} \right|.$$

Let us first consider the choice READ = SUM, so that by (3) we get

$$\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u) \leq \left| \sum_{\alpha=1}^d \theta_\alpha \sum_{i \in \mathbf{V}} \frac{\partial^2 h_i^{(m), \alpha}}{\partial x_u^\beta \partial x_v^\gamma} \right|$$

As before, we index the columns of the Hessian of \mathbf{h}_i as $\frac{\partial^2 h_i^{(m), \alpha}}{\partial x_u^\beta \partial x_v^\gamma} = (\nabla_{uv}^2 \mathbf{h}_i^{(m)})_{\alpha, d(\beta-1)+\gamma}$ and hence obtain

$$\text{mix}_{y_{\mathbf{G}}^{(m)}}(v, u) \leq \sum_{i \in \mathbf{V}} \|(\nabla_{uv}^2 \mathbf{h}_i^{(m)})^\top \boldsymbol{\theta}\| \leq \sum_{i \in \mathbf{V}} \|\nabla_{uv}^2 \mathbf{h}_i^{(m)}\|, \quad (27)$$

since $\boldsymbol{\theta}$ has unit norm. Note that the very same bound in (27) also holds if we replaced the SUM readout with either the MAX or the MEAN readout. We can then rely on Theorem C.2 and find

$$\begin{aligned}
\text{mix}_{y_G^{(m)}}(v, u) &\leq \sum_{i \in \mathcal{V}} \sum_{k=0}^{m-1} \sum_{j \in \mathcal{V}} (c_\sigma \mathbf{w})^{2m-k-1} \mathbf{w}(\mathbf{S}^{m-k})_{jv} (\mathbf{S}^k)_{ij} (\mathbf{S}^{m-k})_{ju} \\
&\quad + c^{(2)} \sum_{i \in \mathcal{V}} \sum_{\ell=0}^{m-1} (c_\sigma \mathbf{w})^{m+\ell} (\mathbf{S}^{m-1-\ell} \mathbf{P}_{(vu)}^{(\ell)})_i \\
&= \sum_{k=0}^{m-1} (c_\sigma \mathbf{w})^{2m-k-1} \left(\mathbf{w}(\mathbf{S}^{m-k})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k) \mathbf{S}^{m-k} \right)_{vu} \\
&\quad + c^{(2)} \sum_{\ell=0}^{m-1} (c_\sigma \mathbf{w})^{m+\ell} (\mathbf{1}^\top \mathbf{S}^{m-1-\ell}) \mathbf{P}_{(vu)}^{(\ell)} \\
&= \sum_{k=0}^{m-1} (c_\sigma \mathbf{w})^{2m-k-1} \left(\mathbf{w}(\mathbf{S}^{m-k})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k) \mathbf{S}^{m-k} \right)_{vu} \\
&\quad + c^{(2)} \sum_{k=0}^{m-1} (c_\sigma \mathbf{w})^{2m-k-1} (\mathbf{1}^\top \mathbf{S}^k) \mathbf{P}_{(vu)}^{(m-k-1)}.
\end{aligned}$$

For the second term we can simply use the formula in (14) to rewrite it in matricial form as claimed – recall that \mathbf{Q}_k is defined in (6). \square

D Proofs and additional details of Section 4

Throughout this Section, for simplicity, we assume $c_\sigma = \max\{|c'_\sigma|, |c''_\sigma|\}$ to be smaller or equal than one – this is satisfied by the vast majority of commonly used non-linear activations, and extending the results below to arbitrary c_σ is straightforward.

D.1 Proof of Theorem 4.2

We begin by proving lower bounds on the spectral norm of the weights, when the depth is the minimal one required to induce any non-zero mixing among nodes v, u . For convenience, we restate Theorem 4.2 here as well.

Theorem 4.2. *Let $\mathbf{A} = \mathbf{A}_{\text{sym}}$, $r := d_G(v, u)$, $m = \lceil r/2 \rceil$, and q be the number of paths of length r between v and u . For an MPNN satisfying Theorem 3.2 with capacity $(m = \lceil r/2 \rceil, \mathbf{w})$, we find $\text{OSQ}_{v,u}(m, \mathbf{w}) \cdot (c_2 \mathbf{w})^r (\mathbf{A}^r)_{vu} \geq 1$. In particular, if the MPNN generates mixing $\text{mix}_{y_G}(v, u)$, then*

$$\mathbf{w} \geq \frac{d_{\min}}{c_2} \left(\frac{\text{mix}_{y_G}(v, u)}{q} \right)^{\frac{1}{r}}.$$

Proof. Without loss of generality, we assume that r is even, so that by our assumptions, we can simply take $m = r/2$. According to Theorem 3.2, we know that the maximal mixing induced by an MPNN of depth m over the features associated with nodes v, u is bounded from above as

$$\text{mix}_{y_G^{(m)}}(v, u) \leq \sum_{k=0}^{m-1} \mathbf{w}^{2m-k-1} \left(\mathbf{w} \mathbf{S}^{m-k} \text{diag}(\mathbf{S}^k \mathbf{1}) \mathbf{S}^{m-k} + c^{(2)} \mathbf{Q}_k \right)_{vu}.$$

where we have replaced c_σ with one, as per our assumption. Since $m = r/2$, where r is the distance among nodes v, u , then the only non-zero contribution for the first term is obtained for $k = 0$ – otherwise we would find a path of length $2(m - k)$ connecting v and u hence violating the assumptions – and is equal to $\mathbf{w}^{2m}(\mathbf{S}^{2m})_{vu}$, and note that $2m = r$. Concerning the terms \mathbf{Q}_k instead, the longest-walk contribution for nodes v, u is $2m - 1$ (when $k = 0$), meaning that $\mathbf{Q}_k = 0$ for all $0 \leq k \leq m - 1$ if $m = 2r$. Accordingly, we can reduce the bound above to:

$$\text{mix}_{y_G^{(m)}}(v, u) \leq \mathbf{w}^{2m} \left(\mathbf{S}^{2m} \right)_{vu} = \mathbf{w}^{2m} \left((c_2 \mathbf{A})^{2m} \right)_{vu},$$

where in the last equality we have again used that $2m = r$, so when expanding the power of \mathbf{S} only the highest-order term in the \mathbf{A} -variable gives non-zero contributions. If we replace now $2m = r$ and use the characterization of over-squashing in Definition 4.1, then

$$\text{OSQ}_{v,u}(m = \frac{r}{2}, \mathbf{w}) \geq (c_2 \mathbf{w})^{-r} \frac{1}{(\mathbf{A}^r)_{vu}}.$$

Therefore, if the MPNN generates mixing $\text{mix}_{y_G^{(m)}}(v, u)$ among the features of v and u , then (8) is satisfied, meaning that the spectral norm of the weights must be larger than

$$\mathbf{w} \geq \frac{1}{c_2} \left(\frac{\text{mix}_{y_G^{(m)}}(v, u)}{(\mathbf{A}^r)_{vu}} \right)^{\frac{1}{r}}.$$

The term \mathbf{A}^r in general can be estimated sharply depending on the knowledge we have of the underlying graph. To get a universal – albeit potentially looser bound – it suffices to note that along each path connecting v and u , the product of the entries of \mathbf{A} can be bounded from above by $(d_{\min})^r$, which completes the proof. \square

We highlight that if $\mathbf{A} = \mathbf{A}_{\text{rw}} = \mathbf{D}^{-1} \mathbf{A}$ – i.e. the aggregation over the neighbours consists of a mean-operation as for the GraphSAGE architecture – then one can apply the very same proof above and derive

Corollary D.1. *The same lower bound for \mathbf{w} in Theorem 4.2, holds when $\mathbf{A} = \mathbf{D}^{-1} \mathbf{A}$.*

D.2 Spectral bounds

Next, we study the case of fixed, bounded spectral norm of the weights, but variable depth, since we are interested in showing that *over-squashing hinders the expressive power of MPNNs for tasks requiring high-mixing of features associated with nodes at high commute time*. We first provide a characterization of the maximal mixing (and hence of the over-squashing measure) in terms of the graph-Laplacian and its pseudo-inverse.

Convention. In the proofs below we usually deal with matrices with nonnegative entries. Accordingly, we introduce the following convention: we write that $\mathbf{A} \leq \mathbf{B}$ if $A_{ij} \leq B_{ij}$ for all entries $1 \leq i, j \leq n$.

Theorem D.2. *Let $\gamma := \sqrt{\frac{d_{\max}}{d_{\min}}}$ and set $\mathbf{A} = \mathbf{A}_{\text{sym}}$ or $\mathbf{A} = \mathbf{A}_{\text{rw}}$. Consider an MPNN as in Thm. 3.2 with depth m , $\max\{\mathbf{w}, \omega/\mathbf{w} + c_1 \gamma + c_2\} \leq 1$. Define $\mathbf{Z} := \mathbf{I} - c_2 \mathbf{\Delta}$. Then the maximal mixing of nodes v, u generated by such MPNN after m layers is*

$$\begin{aligned} \text{mix}_{y_G^{(m)}}(v, u) &\leq \gamma^k \left(m \frac{\sqrt{d_v d_u}}{2|E|} \left(1 + 2c^{(2)}(1 + \gamma^s) \right) + \frac{1}{c_2} \left(\mathbf{Z}^2 (\mathbf{I} - \mathbf{Z}^{2m}) (\mathbf{I} + \mathbf{Z})^{-1} \mathbf{\Delta}^\dagger \right)_{vu} \right) \\ &\quad + 2 \frac{c^{(2)}}{c_2} \gamma^k \left(\left((1 + \gamma^s) \mathbf{I} - \mathbf{\Delta} \right) (\mathbf{I} - \mathbf{Z}^{2m}) (\mathbf{I} + \mathbf{Z})^{-1} \mathbf{\Delta}^\dagger \right)_{vu}, \end{aligned}$$

where $k = s = 1$ if $\mathbf{A} = \mathbf{A}_{\text{sym}}$ or $k = 4, s = 2$ if $\mathbf{A} = \mathbf{A}_{\text{rw}}$.

Proof. We first focus on the symmetrically normalized case $\mathbf{A} = \mathbf{A}_{\text{sym}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, which we recall that we can rewrite as $\mathbf{A}_{\text{sym}} = \mathbf{I} - \mathbf{\Delta}$, where $\mathbf{\Delta}$ is the (normalized) graph Laplacian (9). Since $\mathbf{w} \leq 1$ and the

message-passing matrix is symmetric, by Theorem 3.2, we can bound the maximal mixing of an MPNN as in the statement by

$$\text{mix}_{y_G^{(m)}}(v, u) \leq \sum_{k=0}^{m-1} \left(\mathbf{S}^{m-k} \text{diag}(\mathbf{S}^k \mathbf{1}) \mathbf{S}^{m-k} + c^{(2)} \mathbf{Q}_k \right)_{vu}.$$

We focus on the first sum. Since $\mathbf{A} = \mathbf{D}^{1/2}(\mathbf{D}^{-1}\mathbf{A})\mathbf{D}^{-1/2}$, where $\mathbf{D}^{-1}\mathbf{A}$ is a row-stochastic matrix, we see that

$$\mathbf{S}_{ij} \leq \frac{\omega}{w} \delta_{ij} + c_1 \gamma \delta_{ij} + c_2 \mathbf{A}_{ij},$$

meaning that we can write $\mathbf{S} \leq \alpha \mathbf{I} + c_2 \mathbf{A}$, where $\alpha = \omega/w + c_1 \gamma$, using the convention introduced above. Accordingly, we can estimate the row-sum of the powers of \mathbf{S} by using $(\mathbf{A}^p \mathbf{1})_i \leq \gamma$ as

$$(\mathbf{S}^k \mathbf{1})_i \leq \sum_{p=0}^k \binom{k}{p} \alpha^{k-p} c_2^p (\mathbf{A}^p \mathbf{1})_i \leq \gamma \sum_{p=0}^k \binom{k}{p} \alpha^{k-p} c_2^p = \gamma (\alpha + c_2)^k \leq \gamma,$$

where the last inequality simply follows from the assumptions. Therefore, we find

$$\left(\sum_{k=0}^{m-1} \mathbf{S}^{m-k} \text{diag}(\mathbf{S}^k \mathbf{1}) \mathbf{S}^{m-k} \right)_{vu} \leq \gamma \sum_{k=0}^{m-1} (\mathbf{S}^{2(m-k)})_{vu} = \gamma \sum_{k=1}^m (\mathbf{S}^{2k})_{vu}. \quad (28)$$

By the assumptions on the regularity of the message-functions, we can estimate \mathbf{S} from above by $\mathbf{S} \leq \alpha \mathbf{I} + c_2 \mathbf{A} = (\alpha + c_2) \mathbf{I} - c_2 \mathbf{\Delta} \leq \mathbf{Z}$, and derive

$$\left(\sum_{k=0}^{m-1} \mathbf{S}^{m-k} \text{diag}(\mathbf{S}^k \mathbf{1}) \mathbf{S}^{m-k} \right)_{vu} \leq \gamma \sum_{k=1}^m (\mathbf{Z}^{2k})_{vu}.$$

From the spectral decomposition of the graph-Laplacian in (9) and the properties that $\lambda_0 = 0$ and $\phi_0(v) = \sqrt{d_v/2|E|}$, we find

$$\begin{aligned} \sum_{k=1}^m (\mathbf{Z}^{2k})_{vu} &= \sum_{k=1}^m \sum_{\ell=0}^{n-1} (1 - c_2 \lambda_\ell)^{2k} \phi_\ell(v) \phi_\ell(u) \\ &= m \frac{\sqrt{d_v d_u}}{2|E|} + \sum_{\ell=1}^{n-1} \left(\frac{1 - (1 - c_2 \lambda_\ell)^{2(m+1)}}{1 - (1 - c_2 \lambda_\ell)^2} - 1 \right) \phi_\ell(v) \phi_\ell(u) \\ &= m \frac{\sqrt{d_v d_u}}{2|E|} + \sum_{\ell=1}^{n-1} \left(\frac{(1 - c_2 \lambda_\ell)^2 (1 - (1 - c_2 \lambda_\ell)^{2m})}{(2 - c_2 \lambda_\ell) c_2 \lambda_\ell} \right) \phi_\ell(v) \phi_\ell(u). \end{aligned}$$

Since $c_2 \leq 1$ and \mathbf{G} is not bipartite, we derive that $(\mathbf{I} + \mathbf{Z}) = 2\mathbf{I} - c_2 \mathbf{\Delta}$ is invertible and hence that the following decomposition holds:

$$(\mathbf{I} + \mathbf{Z})^{-1} = \sum_{\ell \geq 0} \frac{1}{2 - c_2 \lambda_\ell} \phi_\ell \phi_\ell^\top.$$

Therefore, we can rely on the spectral-decomposition of the pseudo-inverse of the graph-Laplacian in (10) to get

$$\left(\sum_{k=0}^{m-1} \mathbf{S}^{m-k} \text{diag}(\mathbf{S}^k \mathbf{1}) \mathbf{S}^{m-k} \right)_{vu} \leq \gamma \left(m \frac{\sqrt{d_v d_u}}{2|E|} + \frac{1}{c_2} \left(\mathbf{Z}^2 (\mathbf{I} - \mathbf{Z}^{2m}) (\mathbf{I} + \mathbf{Z})^{-1} \mathbf{\Delta}^\dagger \right)_{vu} \right). \quad (29)$$

It now remains to bound the term $c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu}$. First, we note that by the symmetry of \mathbf{A} and the estimate $(\mathbf{S}^k \mathbf{1})_i \leq \gamma$, that we derived above, we obtain

$$c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} \leq 2c^{(2)} \gamma \left(\sum_{k=0}^{m-1} (\mathbf{A} \mathbf{Z}^{2(m-k-1)})_{vu} + \gamma (\mathbf{Z}^{2(m-k-1)})_{vu} \right).$$

Then we can use the identity $\mathbf{A} = \mathbf{I} - \mathbf{\Delta}$, to find

$$c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} \leq 2c^{(2)}\gamma \sum_{k=0}^{m-1} \left(((1+\gamma)\mathbf{I} - \mathbf{\Delta}) \mathbf{Z}^{2(m-k-1)} \right)_{vu}. \quad (30)$$

By relying on the spectral decomposition as above, we finally get

$$\begin{aligned} c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} &\leq 2c^{(2)}\gamma \left(m \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} (1+\gamma) \right) \\ &\quad + 2c^{(2)}\gamma \left(\sum_{\ell>0} (1+\gamma - \lambda_\ell) \frac{1 - (1 - c_2 \lambda_\ell)^{2m}}{(2 - c_2 \lambda_\ell) c_2 \lambda_\ell} \phi_\ell(v) \phi_\ell(u) \right). \end{aligned}$$

As done previously, we can rewrite the last terms via $(\mathbf{I} + \mathbf{Z})^{-1}$ as

$$\begin{aligned} c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} &\leq 2c^{(2)}\gamma \left(m \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} (1+\gamma) \right) \\ &\quad + 2 \frac{c^{(2)}}{c_2} \gamma \left(((1+\gamma)\mathbf{I} - \mathbf{\Delta})(\mathbf{I} - \mathbf{Z}^{2m})(\mathbf{I} + \mathbf{Z})^{-1} \mathbf{\Delta}^\dagger \right)_{vu}. \end{aligned} \quad (31)$$

We can then combine (29) and (31) and derive the bound we claimed, when $\mathbf{A} = \mathbf{A}_{\text{sym}}$. For the case $\mathbf{A} = \mathbf{A}_{\text{rw}} = \mathbf{D}^{-1}\mathbf{A}$, it suffices to notice that $\mathbf{S} \leq \alpha' \mathbf{I} + c_2 \mathbf{A}$, where $\alpha' = \omega/w + c_1$ and that

$$(\mathbf{1}^\top \mathbf{S}^k)_i \leq \sum_j \sum_{p=0}^k \binom{k}{p} (\alpha')^{k-p} c_2^p ((\mathbf{D}^{-1}\mathbf{A})^p)_{ji} \leq \frac{d_{\max}}{d_{\min}} (\alpha' + c_2)^k \leq \gamma^2,$$

where we have used that by assumption $\alpha' + c_2 \leq 1$. Similarly, we get $(\mathbf{S}^{m-k})^\top \mathbf{S}^{(m-k)} \leq \gamma^2 \mathbf{Z}^{2(m-k)}$. Finally, the \mathbf{Q}_k -term can be bounded by

$$c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} \leq 2c^{(2)} \left(\sum_{k=0}^{m-1} \gamma^4 \mathbf{A}_{\text{sym}} \mathbf{Z}^{2(m-k-1)} + \gamma^6 \mathbf{Z}^{2(m-k-1)} \right),$$

and we can follow the previous steps in the symmetric case to complete the proof. \square

Corollary D.3. *Under the assumptions of Theorem D.2, if the message functions in (2) are linear – as for GCN, SAGE, or GIN – then the maximal mixing induced by such an MPNN of m layers is*

$$\text{mix}_{y_G^{(m)}}(v, u) \leq \left(\frac{d_{\max}}{d_{\min}} \right)^k \left(m \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} + \frac{1}{c_2} \left(\mathbf{Z}^2 (\mathbf{I} - \mathbf{Z}^{2m}) (\mathbf{I} + \mathbf{Z})^{-1} \mathbf{\Delta}^\dagger \right)_{vu} \right)$$

Proof. This follows from Theorem D.2 simply by noticing that if the message-function ψ in (2) is linear, then the upper bound for the norm of the Hessian can be taken to be zero, i.e. $c^{(2)} = 0$. \square

D.3 Proof of Theorem 4.3

We now expand the previous results to derive the minimal number of layers required to induce mixing in the case of bounded weights, showing that the depth may need to grow with the commute time of nodes. We recall that γ is $\sqrt{d_{\max}/d_{\min}}$ while $0 = \lambda_0 < \lambda_1 \leq \dots \leq \lambda_{n-1}$ are the eigenvalues of the symmetrically normalized graph Laplacian (9). We restate Theorem 4.3 below.

Theorem 4.3. Consider an MPNN satisfying Theorem 3.2, with $\max\{\mathbf{w}, \omega/\mathbf{w} + c_1\gamma + c_2\} \leq 1$, and $\mathbf{A} = \mathbf{A}_{\text{sym}}$. If $\text{OSQ}_{v,u}(m, \mathbf{w}) \cdot (\text{mix}_{y_G}(v, u)) \leq 1$, and hence the MPNN generates mixing $\text{mix}_{y_G}(v, u)$ among the features associated with nodes v, u , then the number of layers m satisfies

$$m \geq \frac{\tau(v, u)}{4c_2} + \frac{|\mathbf{E}|}{\sqrt{d_v d_u}} \left(\frac{\text{mix}_{y_G}(v, u)}{\gamma\mu} - \frac{1}{c_2} \left(\frac{\gamma + |1 - c_2\lambda^*|^{r-1}}{\lambda_1} + 2\frac{c^{(2)}}{\mu} \right) \right),$$

where $r = d_G(v, u)$, $\mu = 1 + 2c^{(2)}(1 + \gamma)$ and $|1 - c_2\lambda^*| = \max_{0 < \ell \leq n-1} |1 - c_2\lambda_\ell| < 1$.

Proof. From now on we let r be the shortest-walk distance between v and u . If $m < r/2$, then we incur the under-reaching issue and hence we get zero mixing among the features associated with nodes v, u . Accordingly, we can choose $m \geq r/2$. We need to provide an estimate on the maximal mixing induced by an MPNN as in the statement. We focus on the bound in Theorem 3.2, and recall that the first sum can be bounded as in (28) by

$$\left(\sum_{k=0}^{m-1} \mathbf{S}^{m-k} \text{diag}(\mathbf{S}^k \mathbf{1}) \mathbf{S}^{m-k} \right)_{vu} \leq \gamma \sum_{k=1}^m (\mathbf{S}^{2k})_{vu} \leq \gamma \sum_{k=1}^m (\mathbf{Z}^{2k})_{vu},$$

where $\mathbf{Z} := \mathbf{I} - c_2\Delta$. We can then bound the geometric sum by accounting for the odd powers too. Therefore, we get

$$\gamma \sum_{k=1}^m (\mathbf{Z}^{2k})_{vu} \leq \gamma \sum_{k=0}^{2m} (\mathbf{Z}^k)_{vu} = \gamma \sum_{k=0}^{2m} \sum_{\ell \geq 0} (1 - c_2\lambda_\ell)^k \phi_\ell(v) \phi_\ell(u).$$

As for the proof of Theorem D.2, we separate the contribution of the zero-eigenvalue and that of the positive ones, so we find that

$$\begin{aligned} \sum_{k=0}^{2m} (\mathbf{Z}^k)_{vu} &\leq (2m+1) \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} + \sum_{\ell > 0} \left(\frac{1 - (1 - c_2\lambda_\ell)^{2m+1}}{c_2\lambda_\ell} \right) \phi_\ell(v) \phi_\ell(u) \\ &= (2m+1) \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} + \sum_{\ell > 0} \frac{1}{c_2\lambda_\ell} \phi_\ell(v) \phi_\ell(u) - \sum_{\ell > 0} \frac{(1 - c_2\lambda_\ell)^{2m+1}}{c_2\lambda_\ell} \phi_\ell(v) \phi_\ell(u). \end{aligned} \quad (32)$$

Thanks to the characterization of commute-time provided in (11), we derive

$$\begin{aligned} \sum_{\ell > 0} \frac{1}{c_2\lambda_\ell} \phi_\ell(v) \phi_\ell(u) &= -\frac{\tau(v, u)}{4c_2|\mathbf{E}|} \sqrt{d_v d_u} + \frac{1}{2c_2} \sum_{\ell > 0} \frac{1}{\lambda_\ell} \left(\phi_\ell^2(v) \sqrt{\frac{d_u}{d_v}} + \phi_\ell^2(u) \sqrt{\frac{d_v}{d_u}} \right) \\ &\leq -\frac{\tau(v, u)}{4c_2|\mathbf{E}|} \sqrt{d_v d_u} + \frac{1}{2c_2\lambda_1} \left(\sqrt{\frac{d_v}{d_u}} + \sqrt{\frac{d_u}{d_v}} - \frac{\sqrt{d_v d_u}}{|\mathbf{E}|} \right) \end{aligned} \quad (33)$$

where in the last inequality we have used that $\sum_{\ell > 0} \phi_\ell^2(v) = 1 - \phi_0^2(v)$ since $\{\phi_\ell\}$ constitute an orthonormal basis, with $\phi_0(v) = \sqrt{d_v/2|\mathbf{E}|}$, and that $\lambda_\ell \geq \lambda_1$, for all $\ell > 0$. Next, we estimate the second sum in (32), and we note that λ^* in the statement is either λ_1 or λ_{n-1} :

$$\begin{aligned} -\sum_{\ell > 0} \frac{(1 - c_2\lambda_\ell)^{2m+1}}{c_2\lambda_\ell} \phi_\ell(v) \phi_\ell(u) &\leq \sum_{\ell > 0} \frac{|1 - c_2\lambda^*|^{2m+1}}{c_2\lambda_\ell} |\phi_\ell(v) \phi_\ell(u)| \\ &\leq \frac{|1 - c_2\lambda^*|^{2m+1}}{2c_2\lambda_1} \sum_{\ell > 0} (\phi_\ell^2(v) + \phi_\ell^2(u)) \\ &\leq \frac{|1 - c_2\lambda^*|^r}{2c_2\lambda_1} \left(2 - \frac{d_v}{2|\mathbf{E}|} - \frac{d_u}{2|\mathbf{E}|} \right), \end{aligned} \quad (34)$$

where in the last inequality we have used that $|1 - c_2\lambda^*| < 1$ and that $m \geq r/2$ (otherwise we would have zero-mixing due to under-reaching). Therefore, by combining (33) and (34), we derive that the first sum on the

right hand side of (7) can be bounded from above by

$$\begin{aligned}
\left(\sum_{k=0}^{m-1} \mathbf{S}^{m-k} \text{diag}(\mathbf{S}^k \mathbf{1}) \mathbf{S}^{m-k} \right)_{vu} &\leq \gamma \left((2m+1) \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} - \frac{\tau(v, u)}{4c_2 |\mathbf{E}|} \sqrt{d_v d_u} \right) \\
&\quad + \frac{\gamma}{2c_2 \lambda_1} \left(\sqrt{\frac{d_v}{d_u}} + \sqrt{\frac{d_u}{d_v}} - \frac{\sqrt{d_v d_u}}{|\mathbf{E}|} \right) \\
&\quad + \gamma \frac{|1 - c_2 \lambda^*|^r}{2c_2 \lambda_1} \left(2 - \frac{d_v}{2|\mathbf{E}|} - \frac{d_u}{2|\mathbf{E}|} \right). \tag{35}
\end{aligned}$$

Next, we continue by estimating the second sum entering the right hand side of (7). We recall that by (30), we have

$$\begin{aligned}
c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} &\leq 2c^{(2)} \gamma \sum_{k=0}^{m-1} \left(((1+\gamma)\mathbf{I} - \mathbf{\Delta}) \mathbf{Z}^{2(m-k-1)} \right)_{vu} = 2c^{(2)} \gamma \sum_{k=0}^{m-1} \left(((1+\gamma)\mathbf{I} - \mathbf{\Delta}) \mathbf{Z}^{2k} \right)_{vu} \\
&\leq 2c^{(2)} \gamma \sum_{k=0}^{2(m-1)} \left(((1+\gamma)\mathbf{I} - \mathbf{\Delta}) \mathbf{Z}^k \right)_{vu} \\
&= 2c^{(2)} \gamma \sum_{k=0}^{2(m-1)} \sum_{\ell \geq 0} ((1+\gamma) - \lambda_\ell) (1 - c_2 \lambda_\ell)^k \phi_\ell(v) \phi_\ell(u).
\end{aligned}$$

We then proceed as above, and separate the contributions associated with the kernel of the Laplacian, to find

$$\begin{aligned}
\sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} &\leq 2\gamma \left((1+\gamma)(2m-1) \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} \right) \\
&\quad + 2\gamma(1+\gamma) \sum_{\ell > 0} \frac{1}{c_2 \lambda_\ell} (1 - (1 - c_2 \lambda_\ell)^{2m-1}) \phi_\ell(v) \phi_\ell(u) \tag{36}
\end{aligned}$$

$$-\frac{2\gamma}{c_2} \sum_{\ell > 0} (1 - (1 - c_2 \lambda_\ell)^{2m-1}) \phi_\ell(v) \phi_\ell(u). \tag{37}$$

For the term in (36), we can apply the same estimate as for the case of (32). Similarly, we can bound (37) by

$$-\frac{2\gamma}{c_2} \sum_{\ell > 0} (1 - (1 - c_2 \lambda_\ell)^{2m-1}) \phi_\ell(v) \phi_\ell(u) \leq \frac{2\gamma}{c_2} \sum_{\ell > 0} \frac{1}{2} \left(\phi_\ell^2(v) + \phi_\ell^2(u) \right) \leq \frac{\gamma}{c_2} \left(2 - \frac{d_v}{2|\mathbf{E}|} - \frac{d_u}{2|\mathbf{E}|} \right).$$

Therefore, we can finally bound the \mathbf{Q}_k -terms in (7) by

$$\begin{aligned}
c^{(2)} \sum_{k=0}^{m-1} (\mathbf{Q}_k)_{vu} &\leq 2c^{(2)} \gamma \left((1+\gamma)(2m-1) \frac{\sqrt{d_v d_u}}{2|\mathbf{E}|} - (1+\gamma) \frac{\tau(v, u)}{4c_2 |\mathbf{E}|} \sqrt{d_v d_u} \right) \\
&\quad + c^{(2)} \gamma \left(\frac{1+\gamma}{c_2 \lambda_1} \left(\sqrt{\frac{d_v}{d_u}} + \sqrt{\frac{d_u}{d_v}} - \frac{\sqrt{d_v d_u}}{|\mathbf{E}|} \right) \right) \\
&\quad + c^{(2)} \gamma \frac{1+\gamma}{c_2 \lambda_1} |1 - c_2 \lambda^*|^{r-1} \left(2 - \frac{d_v}{2|\mathbf{E}|} - \frac{d_u}{2|\mathbf{E}|} \right) \\
&\quad + \frac{c^{(2)} \gamma}{c_2} \left(2 - \frac{d_v}{2|\mathbf{E}|} - \frac{d_u}{2|\mathbf{E}|} \right). \tag{38}
\end{aligned}$$

We can the combine (35) and (38), to find that the maximal mixing induced by an MPNN of m layers as in

the statement of Theorem 4.3, is

$$\begin{aligned}
\text{mix}_{y_G}^{(m)}(v, u) &\leq \gamma \sqrt{d_v d_u} \left(\frac{m}{|\mathbf{E}|} \mu + \frac{1}{2|\mathbf{E}|} - \frac{\tau(v, u)}{4c_2 |\mathbf{E}|} \mu \right) \\
&\quad + \gamma \frac{\mu}{2c_2 \lambda_1} \left(\sqrt{\frac{d_v}{d_u}} + \sqrt{\frac{d_u}{d_v}} - \frac{\sqrt{d_v d_u}}{|\mathbf{E}|} \right) \\
&\quad + \gamma \frac{\mu}{2c_2 \lambda_1} |1 - c_2 \lambda^*|^{r-1} \left(2 - \frac{d_v}{2|\mathbf{E}|} - \frac{d_u}{2|\mathbf{E}|} \right) \\
&\quad + \frac{\gamma c^{(2)}}{c_2} \left(2 - \frac{d_v}{2|\mathbf{E}|} - \frac{d_u}{2|\mathbf{E}|} \right), \tag{39}
\end{aligned}$$

where $\mu := 1 + 2c^{(2)}(1 + \gamma)$ and we have removed the term $-2c^{(2)}\gamma(1 + \gamma)\sqrt{d_v d_u}/2|\mathbf{E}| \leq 0$. Moreover, since $\lambda_1 < 1$ unless \mathbf{G} is the complete graph (and if that was the case, then we could take the distance r below to simply be equal to 1) and $c_2 \leq 1$, we find

$$\gamma \sqrt{d_v d_u} \frac{1}{2|\mathbf{E}|} \left(1 - \frac{\mu}{c_2 \lambda_1} \left(1 + \frac{|1 - c_2 \lambda^*|^{r-1}}{2} \left(\sqrt{\frac{d_v}{d_u}} + \sqrt{\frac{d_u}{d_v}} \right) \right) \right) \leq 0.$$

Accordingly, we can simplify (39) as

$$\begin{aligned}
\text{mix}_{y_G}^{(m)}(v, u) &\leq \gamma \sqrt{d_v d_u} \left(\frac{m}{|\mathbf{E}|} \mu - \frac{\tau(v, u)}{4c_2 |\mathbf{E}|} \mu \right) + \frac{\gamma \mu}{2c_2 \lambda_1} \left(\sqrt{\frac{d_v}{d_u}} + \sqrt{\frac{d_u}{d_v}} \right) \\
&\quad + \frac{\gamma \mu}{c_2 \lambda_1} |1 - c_2 \lambda^*|^{r-1} + 2 \frac{\gamma c^{(2)}}{c_2}.
\end{aligned}$$

We can now rearrange the terms and obtain

$$\begin{aligned}
m &\geq \frac{\tau(v, u)}{4c_2} + \frac{|\mathbf{E}|}{\sqrt{d_v d_u}} \left(\frac{\text{mix}_{y_G}(v, u)}{\gamma \mu} - \frac{1}{2c_2 \lambda_1} \left(\sqrt{\frac{d_v}{d_u}} + \sqrt{\frac{d_u}{d_v}} \right) - \frac{1}{c_2 \lambda_1} |1 - c_2 \lambda^*|^{r-1} - 2 \frac{c^{(2)}}{c_2} \frac{1}{\mu} \right) \\
&\geq \frac{\tau(v, u)}{4c_2} + \frac{|\mathbf{E}|}{\sqrt{d_v d_u}} \left(\frac{\text{mix}_{y_G}(v, u)}{\gamma \mu} - \frac{1}{2c_2 \lambda_1} (2\gamma) - \frac{1}{c_2 \lambda_1} |1 - c_2 \lambda^*|^{r-1} - 2 \frac{c^{(2)}}{c_2} \frac{1}{\mu} \right) \\
&\geq \frac{\tau(v, u)}{4c_2} + \frac{|\mathbf{E}|}{\sqrt{d_v d_u}} \left(\frac{\text{mix}_{y_G}(v, u)}{\gamma \mu} - \frac{1}{c_2} \left(\frac{\gamma + |1 - c_2 \lambda^*|^{r-1}}{\lambda_1} + 2 \frac{c^{(2)}}{\mu} \right) \right),
\end{aligned}$$

which completes the proof. \square

We note that the case of $\mathbf{A} = \mathbf{A}_{\text{rw}}$ follows easily since one can adapt the previous argument exactly as in the proof of Theorem D.2, which lead to the same bounds once we replace γ with $\gamma' = d_{\max}/d_{\min}$.

First, we note that the bounds again simplify further and become sharper if the message-functions ψ in (2) are linear.

Corollary D.4. *If the assumptions of Theorem 4.3 are satisfied, and the message-functions ψ are linear – as for GCN, GIN, GraphSAGE – then*

$$m \geq \frac{\tau(v, u)}{4c_2} + \frac{|\mathbf{E}|}{\sqrt{d_v d_u}} \left(\frac{\text{mix}_{y_G}(v, u)}{\gamma} - \frac{1}{c_2 \lambda_1} \left(\gamma + |1 - c_2 \lambda^*|^{r-1} \right) \right).$$

D.4 The case of the unnormalized adjacency matrix

In this Section we extend the analysis on the depth required to induce mixing, to the case of the unnormalized adjacency matrix \mathbf{A} . When $\mathbf{A} = \mathbf{A}$, the aggregation in (2) is simply a sum over the neighbours, a case that

covers the classical GIN-architecture. In this way, the messages are no longer scaled down by the degree of (either) the endpoints of the edge, which means that, *in principle*, the whole GNN architecture is more sensitive but independent of where we are in the graph. First, we generalize Theorem 4.3 to this setting. We note that the same conclusions hold, provided that the maximal spectral norm of the weights is smaller than the maximal degree d_{\max} ; this is not surprising, since it accounts for the lack of the normalization of the messages.

Corollary D.5. *Consider an MPNN as in (2) with $\mathbf{A} = \mathbf{A}$. If $\omega/(\omega d_{\max}) + c_1 + c_2 \leq 1$ and $\omega d_{\max} \leq 1$, then the minimal depth m satisfies the same lower bound as in Theorem 4.3 with $\gamma = 1$.*

Proof. First, we note that in this case

$$\mathbf{S}_{ij} \leq \frac{\omega}{w} \delta_{ij} + c_1 d_{\max} \delta_{ij} + c_2 A_{ij} \leq d_{\max} \left(\alpha \delta_{ij} + c_2 (\mathbf{A}_{\text{sym}})_{ij} \right),$$

where $\alpha = \omega/(\omega d_{\max}) + c_1$. In particular, we find that

$$(\mathbf{S}^k \mathbf{1})_i \leq \sum_{p=0}^k \binom{k}{p} \alpha^{k-p} c_2^p (\mathbf{A}^p \mathbf{1})_i \leq (d_{\max})^k (\alpha + c_2)^k.$$

Accordingly, we can bound the first sum in (7) as

$$\left(\sum_{k=0}^{m-1} w^{2m-k} (d_{\max})^k (\alpha + c_2)^k (d_{\max})^{2(m-k)} \left(\alpha \mathbf{I} + c_2 \mathbf{A}_{\text{sym}} \right)^{2(m-k)} \right)_{vu} \leq \left(\sum_{k=0}^{m-1} \mathbf{Z}^{2(m-k)} \right)_{vu},$$

where we have used the assumptions $\alpha + c_2 \leq 1$, and $\omega d_{\max} \leq 1$, and the definition $\mathbf{Z} := \mathbf{I} - c_2 \mathbf{A}$ in Theorem D.2. Since this term is the same one entering the argument in the proof of Theorem 4.3 (once we set $\gamma = 1$) we can proceed in the same way to estimate it. A similar argument works for the sum of the \mathbf{Q}_k terms, which, thanks to our assumptions, can still be bounded as in (30) with $\gamma = 1$ so that we can finally simply copy the proof of Theorem 4.3. \square

A relative measurement for OSQ. To account for the fact that different message-passing matrices \mathbf{A} may lead to inherently quite distinct scales (think of the case where the aggregation is a mean vs when it is a sum), one could modify the over-squashing characterization in Definition 4.1 as follows:

Definition D.6. *Given an MPNN with capacity (m, w) , we define the **relative over-squashing** of nodes v, u as*

$$\text{OSQ}_{v,u}^{\text{rel}}(m, w) := \left(\frac{\sum_{k=0}^{m-1} w^{2m-k-1} \left(w (\mathbf{S}^{m-k})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k) \mathbf{S}^{m-k} + c^{(2)} \mathbf{Q}_k \right)_{vu}}{\max_{i,j \in \mathcal{V}} \sum_{k=0}^{m-1} w^{2m-k-1} \left(w (\mathbf{S}^{m-k})^\top \text{diag}(\mathbf{1}^\top \mathbf{S}^k) \mathbf{S}^{m-k} + c^{(2)} \mathbf{Q}_k \right)_{ij}} \right)^{-1}.$$

The normalization proposed here is similar to the idea of relative score introduced in [58]. This way, a larger scale induced by a certain choice of the message-passing matrix \mathbf{A} , is naturally accounted for by the relative measurement. In particular, the relative over-squashing is now quantifying the maximal mixing among a certain pair of nodes v, u compared to the maximal mixing that the same MPNN over the same graph can generate among *any* pair of nodes. In our theoretical development in Section 4 we have decided to rely on the absolute measurement since our analysis depends on the derivation of the maximal mixing induced by an MPNN (i.e. upper bounds) which translate into necessary criteria for an MPNN to generate a given level of mixing. In principle, to deal with relative measurements, one would also need some form of lower bound on the maximal mixing and hence address also whether the conditions provided are indeed sufficient. We reserve a thorough investigation of this angle to future work.

E The case of node-level tasks

In this Section we discuss how one can extend our analysis to node-level tasks and further comment on the novelty of our approach compared to existing results in [8, 24]. First, we emphasize that the analysis on the Jacobian of node features carried over in [8, 24] cannot be extended to graph-level functions and that in fact, *our notion of mixing is needed* to assess how two different node-features are communicating when the target is a graph-level function.

From now on, let us consider the case where the function we need to learn is $\mathbf{Y} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$, and as usual we assume it to be equivariant with respect to permutations of the nodes. A natural attempt to connect the results in [8, 24] and the expressivity of the MPNNs – in the spirit of our Section 3 – could be to characterize the **first-order interactions** (or mixing of order 1) of the features associated with nodes v, u with respect to the underlying node-level task \mathbf{Y} as

$$\text{mix}_{\mathbf{Y}}^{(1)}(v, u) = \max_{\mathbf{X}} \max_{1 \leq \alpha, \beta \leq d} \left| \frac{\partial(\mathbf{Y}(\mathbf{X}))_v^\alpha}{\partial x_u^\beta} \right|,$$

where $(\mathbf{Y}(\mathbf{X}))_v \in \mathbb{R}^d$ is the value of the node-level map at v . Accordingly, one can then use Theorem C.1 to derive upper bounds on the maximal first-order interactions that MPNNs (2) can induce among nodes. As a consequence of this approach, we would still find that MPNNs struggle to learn functions with large $\text{mix}_{\mathbf{Y}}^{(1)}(v, u)$ if nodes v, u have large commute time. In particular, in light of Theorem C.1, we can extend the measure of over-squashing to the case of first-order interactions for node-level tasks. Once again, below we tacitly assume that the non-linear activation σ satisfies $|\sigma'| \leq 1$, although it is straightforward to extend the formulation to the general case.

Definition E.1. *Given an MPNN as in (2) with capacity (m, \mathbf{w}) and \mathbf{S} defined in (5), we define the **first-order over-squashing** of v, u as*

$$\text{OSQ}_{v,u}^{(1)}(m, \mathbf{w}) := \left((c_\sigma \mathbf{w})^m (\mathbf{S}^m)_{vu} \right)^{-1}.$$

It follows then from Theorem C.1, that a necessary condition for an MPNN to learn a node-level function \mathbf{Y} with first-order mixing $\text{mix}_{\mathbf{Y}}^{(1)}(v, u)$ is

$$\text{OSQ}_{v,u}^{(1)}(m, \mathbf{w}) < \left(\text{mix}_{\mathbf{Y}}^{(1)}(v, u) \right)^{-1}.$$

It is straightforward to argue as in Theorem 4.3 and [8] for example, to derive that nodes at higher effective resistance will incur higher first-order over-squashing. Accordingly:

An MPNN as in (2) with bounded capacity, cannot learn node-level functions with high first-order interactions among nodes v, u with high effective resistance.

Building a hierarchy of measures. Although first-order derivatives might be enough to capture some form of over-squashing for node-level tasks, even in this scenario we can study the pairwise mixing induced *at a specific node*, and hence consider the curvature (or Hessian) of the node-level function \mathbf{Y} – which is more expressive than the first-order Jacobian. Accordingly, for a node-level function $\mathbf{Y} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$, we say that it has **second-order interactions** (or mixing of order 2) $\text{mix}_{\mathbf{Y}}^{(2)}(i, v, u)$ of the features associated with nodes v, u at a given node i when

$$\text{mix}_{\mathbf{Y}}^{(2)}(i, v, u) = \max_{\mathbf{X}} \left\| \frac{\partial^2(\mathbf{Y}(\mathbf{X}))_i}{\partial \mathbf{x}_u \partial \mathbf{x}_v} \right\|.$$

We can then restate Theorem C.2 as follows – we let $\mathbf{Y}^{(m)}$ be the node-level function computed by an MPNN after m layers.

Corollary E.2. *Given MPNNs as in (2), let σ and $\psi^{(t)}$ be \mathcal{C}^2 functions and assume $|\sigma'|, |\sigma''| \leq c_\sigma$, $\|\boldsymbol{\Omega}^{(t)}\| \leq \omega$, $\|\mathbf{W}^{(t)}\| \leq \mathbf{w}$, $\|\nabla_1 \psi^{(t)}\| \leq c_1$, $\|\nabla_2 \psi^{(t)}\| \leq c_2$, $\|\nabla^2 \psi^{(t)}\| \leq c^{(2)}$. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be defined as in (5). Given nodes*

$i, v, u \in \mathbf{V}$, if $\mathbf{P}_{(vu)}^{(\ell)} \in \mathbb{R}^n$ is as in (14) and m is the number of layers, then the maximal mixing of order 2 of the MPNN at node i satisfies

$$\begin{aligned} \text{mix}_{\mathbf{Y}^{(m)}}^{(2)}(i, v, u) &\leq \sum_{k=0}^{m-1} \sum_{j \in \mathbf{V}} (c_\sigma \mathbf{w})^{2m-k-1} \mathbf{w}(\mathbf{S}^{m-k})_{jv} (\mathbf{S}^k)_{ij} (\mathbf{S}^{m-k})_{ju} \\ &\quad + c^{(2)} \sum_{\ell=0}^{m-1} (c_\sigma \mathbf{w})^{m+\ell} (\mathbf{S}^{m-1-\ell} \mathbf{P}_{(vu)}^{(\ell)})_i. \end{aligned} \quad (40)$$

Similarly to Definition 4.1, we can use the maximal mixing (at the node-level) to characterize the over-squashing of order two at a specific node as follows: as usual, for simplicity we assume that $c_\sigma = 1$.

Definition E.3. Given an MPNN as in (2) with capacity (m, \mathbf{w}) and \mathbf{S} defined in (5), we define the **second-order over-squashing** of v, u at node i as

$$\begin{aligned} \text{OSQ}_{i,v,u}^{(2)}(m, \mathbf{w}) &:= \left(\sum_{k=0}^{m-1} \sum_{j \in \mathbf{V}} \mathbf{w}^{2m-k} (\mathbf{S}^{m-k})_{jv} (\mathbf{S}^k)_{ij} (\mathbf{S}^{m-k})_{ju} \right. \\ &\quad \left. + c^{(2)} \sum_{\ell=0}^{m-1} \mathbf{w}^{m+\ell} (\mathbf{S}^{m-1-\ell} \mathbf{P}_{(vu)}^{(\ell)})_i \right)^{-1}. \end{aligned}$$

It is then straightforward to extend our theoretical analysis to derive how $\text{OSQ}^{(2)}$ prevents MPNNs from learning node-level functions with high-mixing at some specific node i of features associated with nodes v, u at large commute time. To support our claim, consider the setting in Theorem 4.3 and hence let $\mathbf{A} = \mathbf{A}_{\text{sym}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$. Under the same assumptions of Theorem 4.3, we find

$$(\mathbf{S}^k)_{ij} \leq 1.$$

We can then simply copy the proof of Theorem 4.3 once we set $\gamma = 1$ and extend its conclusions as follows:

Corollary E.4. Consider an MPNN as in (2) and let the assumptions of Theorem 4.3 hold. If the MPNN generates second-order mixing $\text{mix}_{\mathbf{Y}^{(2)}}^{(2)}(i, v, u)$ at node i , with respect to the features associated with nodes v, u , then the number of layers m satisfy:

$$m \geq \frac{\tau(v, u)}{4c_2} + \frac{|E|}{\sqrt{d_v d_u}} \left(\frac{\text{mix}_{y_G}(v, u)}{\mu} - \frac{1}{c_2} \left(\frac{1 + |1 - c_2 \lambda^*|^{r-1}}{\lambda_1} + 2 \frac{c^{(2)}}{\mu} \right) \right),$$

where $\mu = 1 + 4c^{(2)}$.

Accordingly, in this Section we have adapted our results from graph-level tasks to node-level tasks and proved that:

The message of Section E. An MPNN of bounded capacity (m, \mathbf{w}) , cannot learn node-level functions that, at some node i , induce high (first order or second order) mixing of features associated with nodes v, u whose commute time is large.

F Additional details of experiments and further ablations 5

F.1 Computing the commute time

The commute time τ between two nodes $u, v \in \mathbf{V}$ on a graph \mathbf{G} can be efficiently computed via the effective resistance \mathbf{R} , with $\tau(u, v) = 2|E|\mathbf{R}(u, v)$. In order to compute the effective resistance \mathbf{R} , we introduce the

(non-normalized) Laplacian matrix $\mathbf{L} = \mathbf{D} - \mathbf{A}$, where \mathbf{D} is the degree matrix. The effective resistance can then be computed by

$$R(u, v) = \Gamma_{uu} + \Gamma_{vv} - 2\Gamma_{uv},$$

where Γ is the the Moore-Penrose inverse of

$$\mathbf{L} + \frac{1}{|V|} \mathbf{1}_{|V| \times |V|},$$

with $\mathbf{1}_{|V| \times |V|} \in \mathbb{R}^{|V| \times |V|}$ being a matrix with all entries set to one.

F.2 Statistics of the results

Table 2: Mean and standard deviation (multiplied by a factor of 10^2) of the results reported in Fig. 3.

Model	$\alpha = 0.0$	$\alpha = 0.2$	$\alpha = 0.4$	$\alpha = 0.6$	$\alpha = 0.8$	$\alpha = 1.0$
GCN	0.80 ± 0.027	1.92 ± 0.039	1.89 ± 0.037	2.21 ± 0.030	3.26 ± 0.061	6.02 ± 0.012
GIN	0.34 ± 0.014	0.68 ± 0.183	1.08 ± 0.353	1.82 ± 0.291	2.78 ± 0.140	5.48 ± 0.683
GraphSAGE	0.04 ± 0.002	0.30 ± 0.020	0.58 ± 0.032	0.81 ± 0.034	1.39 ± 0.018	6.43 ± 0.017
GatedGCN	0.01 ± 0.001	0.06 ± 0.002	0.14 ± 0.007	0.27 ± 0.014	0.46 ± 0.014	6.46 ± 0.011

Table 3: Mean and standard deviation (multiplied by a factor of 10^2) of the results reported in Fig. 4.

Model	$m = 5$	$m = 10$	$m = 15$	$m = 20$	$m = 25$	$m = 32$
GCN	6.12 ± 0.003	3.53 ± 0.017	2.13 ± 0.003	1.55 ± 0.048	1.50 ± 0.392	1.14 ± 0.024
GIN	6.06 ± 0.010	1.74 ± 0.015	1.48 ± 0.424	1.89 ± 0.010	1.05 ± 0.123	1.02 ± 0.093
GraphSAGE	6.40 ± 0.006	1.84 ± 0.068	0.55 ± 0.041	0.61 ± 0.007	0.46 ± 0.035	0.44 ± 0.119
GatedGCN	6.40 ± 0.001	0.39 ± 0.049	0.29 ± 0.025	0.35 ± 0.032	0.43 ± 0.045	0.82 ± 0.171

F.3 On the training error

In this section, we report the training error of the MPNNs trained in section 5. Fig. 5 shows the training MAE corresponding to the experiment in section 5.1, while Fig. 6 shows the training MAE corresponding to section 5.2. We can see that in both cases, the training MAE exhibits the same qualitative behavior as the reported test MAE in the main paper, i.e., the training MAE increases for increasing levels of commute time τ , while it decreases for increasing number of MPNN layers, which further validates our claim that *over-squashing hinders the expressive power of MPNNs*.

F.4 On the over-squashing measure

In this section, we examine how the over-squashing measure OSQ (as of Definition 4.1) depends on the commute time τ as well as on the depth of the underlying MPNN. To this end, we follow the experimental setup of section 5.1 and 5.2, but instead of training the models and presenting their performance in terms of the test MAE, we compute OSQ of the underlying models. We can see in Fig. 7 that OSQ increases for increasing values of the α -quantile of the τ -distribution for all MPNNs considered here. Moreover, we can see in Fig. 8 that OSQ decreases for increasing number of layers for all considered models.

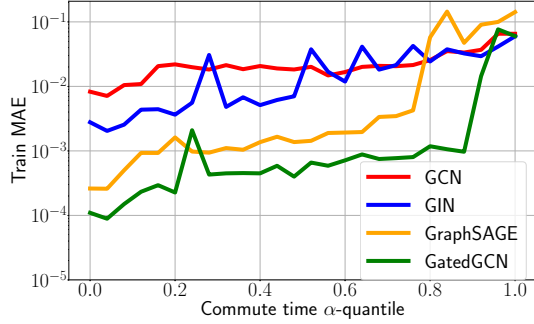


Figure 5: Train MAE of GCN, GIN, GraphSAGE, and GatedGCN on synthetic ZINC, where the commute time of the underlying mixing is varied, while the MPNN architecture is fixed (e.g., depth, number of parameters), i.e., mixing according to increasing values of the α -quantile of the τ -distribution over the ZINC graphs.




Figure 6: Train MAE of GCN, GIN, GraphSAGE, and GatedGCN on synthetic ZINC, where the commute time is fixed to be high (i.e., at the level of the 0.8-quantile), while only the depth of the underlying MPNN is varied between 5 and 32 (all other architectural components are fixed).




Figure 7: OSQ (Definition 4.1) of GCN, GIN, GraphSAGE, and GatedGCN on synthetic ZINC, where the commute time of the underlying mixing is varied, while the MPNN architecture is fixed (e.g., depth, number of parameters), i.e., mixing according to increasing values of the α -quantile of the τ -distribution over the ZINC graphs.




Figure 8: OSQ (Definition 4.1) of GCN, GIN, GraphSAGE, and GatedGCN on synthetic ZINC, where the commute time is fixed to be high (i.e., at the level of the 0.8-quantile), while only the depth of the underlying MPNN is varied between 5 and 32 (all other architectural components are fixed).