
Pruning at Initialisation through the lens of Graphon Limit: Convergence, Expressivity, and Generalisation

Hoang Pham¹ The-Anh Ta² Long Tran-Thanh¹

Abstract

Pruning at Initialisation methods discover sparse, trainable subnetworks before training, but their theoretical mechanisms remain elusive. Existing analyses are often limited to finite-width statistics, lacking a rigorous characterisation of the global sparsity patterns that emerge as networks grow large. In this work, we connect discrete pruning heuristics to graph limit theory via graphons, establishing the *graphon limit of PaI masks*. We introduce a *Factorised Saliency Model* that encompasses popular pruning criteria and prove that, under regularity conditions, the discrete masks generated by these algorithms converge to deterministic bipartite graphons. This limit framework establishes a novel topological taxonomy for sparse networks: while unstructured methods (e.g., Random, Magnitude) converge to homogeneous graphons representing uniform connectivity, data-driven methods (e.g., SNIP, GraSP) converge asymptotically to heterogeneous graphons that encode implicit feature selection. Leveraging this continuous characterisation, we derive two consequences. First, we prove a universal approximation theorem for sparse networks on active coordinate subspaces. Second, under the Graphon-NTK lazy-training regime, we connect the limiting graphon to NTK-style generalisation bounds and introduce a path-density interpretation of how sparse topology can modulate kernel alignment. Our results transform the study of sparse neural networks from combinatorial graph problems into a rigorous framework of continuous operators, offering a new mechanism for analysing expressivity and generalisation in sparse networks.

¹Department of Computer Science, University of Warwick ²CSIRO's Data61. Correspondence to: Hoang Pham <hoang.pham@warwick.ac.uk>.

1. Introduction

Pruning at Initialisation (PaI) has established itself as a powerful paradigm for efficient deep learning, enabling the discovery of sparse, trainable subnetworks without the prohibitive cost of dense pre-training (Lee et al., 2019b; Wang et al., 2020; Tanaka et al., 2020; Liu et al., 2022; Pham et al., 2023; Xiang et al., 2025). Inspired by the Lottery Ticket Hypothesis (LTH) (Frankle & Carbin, 2018; Frankle et al., 2020; Zhou et al., 2019), PaI methods such as SNIP (Lee et al., 2019b), GraSP (Wang et al., 2020), and SynFlow (Tanaka et al., 2020) utilise initialisation data to identify pruning masks for producing sparse subnetworks whose performance often matches dense networks. Despite their empirical success, the theoretical nature of the resulting sparse architectures remains opaque (Pham et al., 2025). Prior works have made brilliant strides in analysing global sparse structure using finite-width combinatorial metrics, such as effective path and node balancing (Pham et al., 2023), path kernels (Gebhart et al., 2021; Patil & Dvornik, 2021), and graph-theoretic connectivity properties (Hoang et al., 2023). However, translating these discrete, finite-width heuristics into a rigorous continuum limit remains an open challenge. A fundamental question is whether the discrete, algorithmic masks produced by these methods converge to stable mathematical objects as network width grows, or whether they remain fundamentally chaotic.

Understanding this asymptotic behaviour is necessary for developing a rigorous theory of sparse deep learning. While classical results rely on fixed random structures (e.g., Erdős-Rényi graphs) (Yang & Wang, 2023), modern PaI methods induce data-dependent correlations between input features and network weights. Recently, (Pham et al., 2025) proposed the Graphon Limit Hypothesis (GLH), postulating that pruning masks converge to continuous graphon operators, which are limit objects of convergent graph sequences (Lovász & Szegedy, 2006; Borgs et al., 2008). In particular, graphon provides a natural framework for describing large-scale network topology. Then, Pham et al. (2025) leveraged this graphon existence assumption to analyse optimisation dynamics via a Graphon-NTK - a sparse network counterpart of the fundamental Neural Tangent Kernel (NTK) of dense models (Jacot et al., 2018; Arora et al., 2019b; Lee et al., 2019a). While this is already a significant conceptual

leap, this hypothesis remains an assumption rather than a derived consequence of the algorithms themselves. A rigorous characterisation of which specific graphons arise from pruning criteria, and how these limit objects govern the network’s expressivity and generalisation capabilities beyond optimisation, is still missing.

In this paper, we fill this gap by providing the first rigorous derivation of the graphon limits for a broad class of popular PaI methods. By modelling saliency scores via a newly proposed Factorised Saliency Model, we prove that under standard infinite-width conditions, the discrete masks generated by these PaI algorithms converge in probability to deterministic bipartite graphons. This result validates the GLH (Pham et al., 2025) and provides a mathematical foundation for treating sparse networks as discretisations of continuous operators.

Our analysis reveals that the limit graphon acts as a distinct asymptotic signature (in cut metric) for pruning algorithms, distinguishing methods whose masks appear similar at finite widths. We show that methods like Random and Magnitude converge to *homogeneous graphons* (constant functions). This implies that in the infinite-width limit, these methods essentially perform uniform downscaling weight matrices, which does not induce data-dependent structure beyond uniform sparsity. In contrast, gradient-based PaI methods like SNIP, or GraSP converge asymptotically to *heterogeneous graphons*. Specifically, the connectivity probability between layers is governed by input features $\varphi(u)$ and output neuron sensitivity $\psi(v)$. This mathematically formalises how data-dependent pruning performs implicit feature selection, concentrating connectivity density on “active” coordinates while filtering out noise. By lifting the analysis from discrete graphs to continuous graphons, we derive properties that are difficult to see in the finite regime. To summarise, our main contributions are:

- **Rigorous convergence analysis of PaI masks towards graphon limit:** We prove that masks generated by factorised saliency scores converge in the cut-norm metric to a deterministic limit graphon (Theorem 4.7). We further show how methods like SNIP and GraSP admit asymptotically rank-equivalent approximations that satisfy this model, formally justifying the GLH.
- **Universal approximation theorem for sparse networks:** We address the expressivity of sparse networks in the high-dimensional limit. We prove that networks sparsified by graphons retain the universal approximation capabilities of dense networks, provided the target function lies within the “active” coordinate subspace discovered by the saliency scores. By concentrating the sparsity budget on relevant features, data-driven pruning inherently preserves the necessary dense computational cores that Random pruning destroys at extreme sparsities.
- **Graphon-modulated generalisation interpretation.** Building on the Graphon-NTK framework (Pham et al., 2025), we show how the limiting graphon enters standard NTK-style generalisation bounds (Cao & Gu, 2019; Arora et al., 2019a) through the induced kernel complexity term. We further introduce a path-density proxy that interprets how sparse topology redistributes input-output connectivity. This provides a mechanistic explanation, supported by empirical evidence, for why data-dependent pruning can outperform uniform Random pruning at extreme sparsity.

2. Related works

Sparse neural networks. The observation that dense networks contain sparse, trainable subnetworks was popularised by the Lottery Ticket Hypothesis (Frankle & Carbin, 2018; Frankle et al., 2020; Burkholz, 2022a;b). To identify these subnetworks without expensive pre-training, Pruning at Initialisation (PaI) approach is getting attention. In particular, various “single-shot” criteria have been proposed, including Magnitude Pruning (Han et al., 2015), which selects weights based on absolute value, while data-driven methods like SNIP (Lee et al., 2019b), GraSP (Wang et al., 2020) utilise gradient information to estimate connection sensitivity. Other PaI methods, such as (Tanaka et al., 2020; Liu et al., 2024), iteratively prune the network until the targeted sparsity is achieved. Later, (Frankle et al., 2021; Su et al., 2020) empirically observe that applying layer-wise connection shuffling or weight reinitialisation on subnetworks found by those PaI methods retains subnetworks’ performance. This raises questions about what the actual aspects are in finding lottery tickets. Prior works analyse the global structure using finite-width combinatorial metrics, including discrete path/node counting (Pham et al., 2023), path kernel (Gebhart et al., 2021; Patil & Dovrolis, 2021), and graph-theoretic properties (Hoang et al., 2023). Our novelty lies in transitioning to rigorous continuous limits through the graphon object. By characterising the asymptotic graphon limit of these masks, we provide a theoretical tool to analyse the effectiveness of sparse networks.

Neural tangent kernels and generalisation error bounds. The behaviour of neural networks as widths $n \rightarrow \infty$ is typically analysed through the lens of the NTK (Jacot et al., 2018; Arora et al., 2019b; Lee et al., 2019a;a). In this regime, under specific parameterisations, the network evolves as a linear model over a fixed kernel Θ_∞ . Extensive work has established generalisation bounds for this regime. Arora et al. (2019a); Cao & Gu (2019) derived rigorous bounds scaling as $O(\sqrt{y^\top \Theta_\infty^{-1} y/m})$ where m is the number of training samples. Crucially, these bounds depend on the Reproducing Kernel Hilbert Space (RKHS) norm of the target function, implying that generalisation is governed by the alignment between the fixed kernel and the target task. While powerful, standard NTK analysis typically assumes

dense connectivity.

Graphon limit of sparse neural networks. Recently, Pham et al. (2025) postulated the Graphon Limit Hypothesis and defined the "Graphon-NTK" to analyse the optimisation dynamics of pruned networks. In particular, Pham et al. (2025) assumed that pruning masks converge to graphons when the network's widths tend to infinity. The concept of graphons, limit objects describing the limits of convergent sequences of graphs, is used to analyse the behaviour of large graphs (Lovász & Szegedy, 2006; Borgs et al., 2008). In the machine learning domain, graphons have primarily been applied to Graph Neural Networks (GNNs) (Ruiz et al., 2020; Keriven et al., 2020; Krishnagopal & Ruiz, 2023), which utilised graphons to define convolution operations that are stable across graphs of varying sizes, effectively enabling "transfer learning" for GNNs on large-scale topological data. Our application differs fundamentally. Rather than using graphons to model input data or statistical networks, we employ them to model the internal architecture of the neural network itself. While Pham et al. (2025) recently introduced this perspective to analyse training dynamics, we deepen this connection by proving that the specific bipartite structures induced by PaI algorithms correspond to known classes of graphons.

Approximation theory for neural networks. Classical universal approximation theorems, which show that sufficiently large dense networks can well approximate any continuous function on a bounded domain, are well-established (Cybenko, 1989; Hornik, 1991; Barron, 1994; Eldan & Shamir, 2016; Nakada & Imaizumi, 2020). In contrast, the theoretical capability of sparse networks to approximate functions has been studied primarily through two lenses: constructive existence and random sparsity. Bolcskei et al. (2019) established fundamental lower bounds on the connectivity required to approximate specific function classes, often relating optimal sparse architectures to affine systems like wavelets and shearlets. While these results prove that optimal sparse networks exist, they do not guarantee that standard pruning algorithms can find them. Conversely, the "Strong Lottery Ticket Hypothesis" literature led by Malach et al. (2020); Orseau et al. (2020), focuses on the probability that a sufficiently wide random network contains a good sparse subnetwork that, without training, still approximates the target function with high accuracy. Instead of assuming a generic random mask or a handcrafted affine system, we derive the approximation specifically for the sparse networks induced by limit graphons.

3. Preliminaries and settings

3.1. Preliminaries

Graphons provide a continuous representation of large discrete graphs (Lovász & Szegedy, 2006; Borgs et al., 2008). For bipartite networks connecting input coordi-

nates to hidden neurons, we work with measurable kernels $\mathcal{W} : [0, 1]^2 \rightarrow [0, 1]$ that encode the connectivity pattern between an infinite set of nodes (Pham et al., 2025).

Bipartite graphons and step kernels. Given a matrix $A \in [0, 1]^{d \times n}$, we define its associated *step kernel* $W_A : [0, 1]^2 \rightarrow [0, 1]$ by $W_A(u, v) := A_{ij}$ for $u \in I_i^{(d)} := \left(\frac{i-1}{d}, \frac{i}{d}\right]$, $v \in J_j^{(n)} := \left(\frac{j-1}{n}, \frac{j}{n}\right]$. This construction provides the standard embedding of bipartite adjacency matrices into the space of measurable kernels on $[0, 1]^2$.

Cut norm and cut distance. The cut norm measures the maximum discrepancy in average mass over all measurable rectangles. For an integrable kernel $\mathcal{W} : [0, 1]^2 \rightarrow [0, 1]$, the *cut norm* is defined as

$$\|\mathcal{W}\|_{\square} := \sup_{S, T \subseteq [0, 1] \text{ measurable}} \left| \int_{S \times T} \mathcal{W}(u, v) du dv \right|. \quad (1)$$

The *bipartite cut distance* between kernels \mathcal{W} and \mathcal{U} is

$$\delta_{\square}^{\text{bip}}(\mathcal{W}, \mathcal{U}) := \inf_{\phi, \psi} \|\mathcal{W} - \mathcal{U}^{\phi, \psi}\|_{\square} \quad (2)$$

where $\mathcal{U}^{\phi, \psi}(u, v) := \mathcal{U}(\phi(u), \psi(v))$ and the infimum is over all measure-preserving maps $\phi, \psi : [0, 1] \rightarrow [0, 1]$. Crucially, in the bipartite setting, row and column relabelings are applied *independently*, unlike the symmetric graphon setting. Hence, we may sort rows/columns by latent features to reveal structure.

3.2. Setting and notations

We study pruning-at-initialisation in a one-hidden-layer network with hidden width n on input $x \in \mathbb{R}^d$

$$f_n(x) = \frac{1}{\sqrt{n}} \sum_{j=1}^n \theta_j^{(2)} \sigma(h_j^{(1)}), \quad h_j^{(1)} := \frac{1}{\sqrt{d}} \sum_{i=1}^d \theta_{ij}^{(1)} x_i,$$

and squared loss on a fixed label $y \in \mathbb{R}$: $\mathcal{L} = \frac{1}{2} (f_n(x) - y)^2$, with i.i.d. Gaussian initialization $\theta_{ij}^{(1)} \sim \mathcal{N}(0, 1)$, i.i.d. output weights $\theta_j^{(2)} \sim \mathcal{N}(0, 1)$, independent of $\theta_{ij}^{(1)}$. Here, the first layer weight matrix is $\theta^{(1)} \in \mathbb{R}^{d \times n}$. We focus on *PaI sparsification of the first layer*: a binary mask $M_n \in \{0, 1\}^{d \times n}$ is computed at initialisation and applied elementwise to $\theta^{(1)}$ as $\theta^{(1)} \leftarrow \theta^{(1)} \odot M_n$. After this, training proceeds only on the active parameters. Our theoretical results concern the asymptotic structure of the mask M_n as $d, n \rightarrow \infty$.

Masks as bipartite graphs and step-kernel embedding.

The mask M_n can be viewed as the adjacency matrix of a bipartite graph between input coordinates and hidden units. To study limits as $d, n \rightarrow \infty$, we associate M_n with a step kernel $W_{M_n} : [0, 1]^2 \rightarrow \{0, 1\}$ defined by $W_{M_n}(u, v) = (M_n)_{ij}$ for $u \in \left(\frac{i-1}{d}, \frac{i}{d}\right]$, $v \in \left(\frac{j-1}{n}, \frac{j}{n}\right]$. We measure

convergence of these step kernels using the *bipartite cut distance*. From now on, we use W_n interchangeably with W_{M_n} for brevity.

Other notations. We write $[k] := \{1, \dots, k\}$ for any $k \in \mathbb{N}^+$. $A_{n,ij}$ or $(A_n)_{ij}$ denotes the (i, j) entry of size n matrix, and \odot denotes Hadamard multiplication.

4. Graphon convergence of PaI methods

Pruning-at-Initialisation methods output *binary mask matrices* $M_n \in \{0, 1\}^{d \times n}$ indicating which connections survive. However, analysing M_n directly is non-trivial due to (i) the mask is a high-dimensional combinatorial object, and (ii) reordering neurons produces different masks with identical functionality. Graphon theory provides a principled way to talk about what such masks look like at the large-width limit. Concretely, we embed M_n into a step kernel W_{M_n} on $[0, 1]^2$ and measure convergence in the (bi)partite cut distance, which compares two kernels *after optimising over all measure-preserving transformations*. In this section, we ask whether the mask has a stable *large-scale connectivity pattern* up to relabeling, i.e. whether W_{M_n} converges in cut distance to a deterministic graphon \mathcal{W} .

4.1. The factorised saliency model

Our analysis is driven by a simple structural observation shared by many PaI criteria where edge saliency often decomposes into an *input feature* (forward signal), a *neuron feature* (backward/gradient signal), and an *edge-level* randomness from initialisation. We formalise this via the factorised saliency model:

$$(S_n)_{ij} = \varphi_{n,i} \cdot \psi_{n,j} \cdot |\xi_{n,ij}|, \quad (3)$$

$$(M_n)_{ij} = \mathbf{1}\{(S_n)_{ij} > \tau_n\}, \quad i \in [d], j \in [n], \quad (4)$$

where $\tau_n > 0$ is a pruning threshold; $\varphi_{n,i} \geq 0$ is *input feature*, capturing forward-pass saliency of input coordinate i ; $\psi_{n,j} \geq 0$ is *neuron feature*, capturing backward-pass (gradient) saliency of hidden unit j ; $\xi_{n,ij}$ is *edge noise*, capturing initialisation randomness at edge (i, j) . This model helps us state and prove the graphon convergence theorem, which can apply to various PaI methods.

PaI instantiations. Many popular PaI methods fit this model either exactly (e.g., Magnitude/Random at i.i.d. initialisation), or asymptotically ranking-equivalent (e.g., SNIP/GraSP in the infinite width regimes) in the one-hidden-layer network setting. Table 1 summarises the corresponding (φ, ψ, ξ) choices. Please refer to Appendix B, C for detailed derivations.

Magnitude/Random. Magnitude pruning sets $(S_n)_{ij} = |\theta_{ij}^{(1)}|$, thus $\varphi_i \equiv \psi_j \equiv 1$. Random pruning (Bernoulli(ρ)) corresponds to $\varphi_i \equiv \psi_j \equiv 1$ with $\xi_{ij} \sim \text{Unif}[0, 1]$.

SNIP (Lee et al., 2019b). With $f_n(x) = \frac{1}{\sqrt{n}} \sum_{j=1}^n \theta_j^{(2)} a_j^{(1)}$ and $h_j^{(1)} = \frac{1}{\sqrt{d}} \sum_{i=1}^d \theta_{ij}^{(1)} x_i$, SNIP introduces masks m_{ij} via $\theta_{ij}^{(1)} \mapsto m_{ij} \theta_{ij}^{(1)}$ and evaluates at $m \equiv \mathbf{1}$. By the chain rule,

$$S_{ij}^{\text{SNIP}} = \left| \frac{\partial \mathcal{L}}{\partial m_{ij}} \right| = \frac{|\delta|}{\sqrt{nd}} |x_i| (|\theta_j^{(2)}| |\sigma'(h_j^{(1)})|) |\theta_{ij}^{(1)}|, \quad (5)$$

where δ denotes the scalar loss derivative w.r.t. the network output (e.g., $\delta = f_n(x) - y$ for squared loss). Absorbing the global factor $|\delta|/\sqrt{nd}$ into τ_n yields the scale-free factorisation in Table 1.

GraSP (Wang et al., 2020) (*magnitude variant*). GraSP assigns a score proportional to $\theta^{(1)} \odot (Hg)$ where $g = \nabla_{\theta^{(1)}} \mathcal{L}(\theta^{(1)})$ and $H = \nabla_{\theta^{(1)}}^2 \mathcal{L}(\theta^{(1)})$. To align with our non-negative thresholding Eqn.4, we use the magnitude variant $S_{ij}^{\text{GraSP}} := |\theta_{ij}^{(1)} (Hg)_{ij}|^1$. In the one-hidden-layer setting, we have $g_{ij} = \partial \mathcal{L} / \partial \theta_{ij}^{(1)} = \delta \frac{1}{\sqrt{n}} \theta_j^{(2)} \sigma'(h_j^{(1)}) \frac{1}{\sqrt{d}} x_i$. For squared loss, $H = J^\top J + \delta \nabla_{\theta^{(1)}}^2 f_n$ with $J = \nabla_{\theta^{(1)}} f_n$. A direct computation yields $(J^\top J g)_{ij} = c_n g_{ij}$ where $c_n := \frac{\|x\|_2^2}{d} \frac{1}{n} \sum_{k=1}^n (\theta_k^{(2)})^2 \sigma'(h_k^{(1)})^2$. Under infinite width setting, the term $\delta (\nabla_{\theta^{(1)}}^2 f_n) g$ contributes only lower-order corrections, hence $(Hg)_{ij} = c_n g_{ij} + o(1)$ entrywise. Thus, up to a global positive factor (irrelevant for top- ρ ranking), $S_{ij}^{\text{GraSP}} \approx \varphi_i^{(n)} \psi_j^{(n)} |\xi_{ij}^{(n)}| = |x_i| (|\theta_j^{(2)}| |\sigma'(h_j^{(1)})|) |\theta_{ij}^{(1)}|$. (6)

SynFlow (Tanaka et al., 2020) (*one iteration*). SynFlow uses the data-agnostic synaptic saliency $S(\theta) = (\partial R / \partial \theta) \odot \theta$ with $R = \mathbf{1}^\top (\prod_{\ell=1}^L |\theta^{(\ell)}|) \mathbf{1}$. In a single-hidden-layer network, one iteration yields $S_{ij} \propto |\theta_j^{(2)}| |\theta_{ij}^{(1)}|$, corresponding to $\varphi_i \equiv 1$ and $\psi_j = |\theta_j^{(2)}|$.

4.2. Regularity assumptions

We establish convergence as $d \rightarrow \infty$ and $n \rightarrow \infty$ under the following conditions:

Assumption 4.1 (Asymptotic width scaling). $\frac{\log(d+n)}{\min\{d,n\}} \rightarrow 0$.

Assumption 4.2 (Deterministic input features). There exists a bounded continuous function $\varphi : [0, 1] \rightarrow [0, \infty)$ such that, after a permissible relabeling of input coordinates, $\max_{1 \leq i \leq d} |\varphi_{n,i} - \varphi(i/d)| \xrightarrow{\mathbb{P}} 0$.

Assumption 4.3 (Neuron features empirical CDF convergence). Let $\widehat{F}_{\psi,n}(t) := \frac{1}{n} \sum_{j=1}^n \mathbf{1}\{\psi_{n,j} \leq t\}$. There exists a deterministic CDF F_ψ supported on $[0, \infty)$ such that

$$\sup_{t \in \mathbb{R}} |\widehat{F}_{\psi,n}(t) - F_\psi(t)| \xrightarrow{\mathbb{P}} 0.$$

Assumption 4.4 (Conditional independence of edge noise). Let $\mathcal{G}_n := \sigma(\{\varphi_{n,i}\}_{i \leq d}, \{\psi_{n,j}\}_{j \leq n}, \tau_n)$. Conditional on

¹We provide analysis for original version in Appendix C

\mathcal{G}_n , the noises $\{\xi_{n,ij}\}_{i \leq d, j \leq n}$ are independent with common law P_ξ , and are independent of \mathcal{G}_n . Moreover, the CDF of $|\xi|$ on $[0, \infty)$ is continuous.

Assumption 4.5 (Deterministic threshold). There exists $\tau \in (0, \infty)$ such that $\tau_n \xrightarrow{P} \tau$, and τ_n is \mathcal{G}_n -measurable.

Remark 4.6. Our aim is to show that the masked weight matrix M_n behaves, at large (d, n) , like a bipartite random graph generated from a *deterministic* graphon $\mathcal{W}(u, v)$: the row index i becomes a continuum coordinate $u = i/d$, the column index j becomes $v = j/n$, and edges are kept with a probability that depends smoothly on the product of a row feature and a column feature. Assumptions 4.2 - 4.3 provide exactly these stable continuum coordinates by requiring that (after relabelling) the input (row) features $\{\varphi_{n,i}\}$ track a deterministic function $\varphi(u)$ and the neuron (column) features $\{\psi_{n,j}\}$ have a deterministic limiting distribution (hence order statistics $\psi_{n,(j)} \approx Q_\psi(v)$, where $Q_\psi(v) := \inf\{y \geq 0 : F_\psi(y) \geq v\}$ is the generalised quantile). Assumption 4.4 isolates the remaining randomness to i.i.d. edge noise $\{\xi_{n,ij}\}$ conditional on $(\varphi_n, \psi_n, \tau_n)$, which turns the mask into a conditionally independent Bernoulli matrix and enables cut-norm concentration around its conditional mean. Assumption 4.5 says that the top- ρ threshold stabilises to a constant τ , so the limiting edge-probability kernel is well-defined and does not inherit extra global randomness. Finally, assumption 4.1 is a mild growth condition ensuring that these concentration bounds hold *uniformly*, so the full bipartite cut distance converges.

SNIP is asymptotically equivalent to a factorised model.

In the 1-hidden layer SNIP setting in Eqn.5, the factorisation Eqn.3 holds after absorbing the global factor into τ_n , with $\varphi_{n,i} = |x_i|$, $\psi_{n,j} = |\theta_j^{(2)}| |\sigma'(h_j^{(1)})|$, and $\xi_{n,ij} = \theta_{ij}^{(1)}$. If the input coordinates are i.i.d. and bounded, then after sorting $\{|x_i|\}_{i \leq d}$ one has $\max_i ||x_{(i)}| - Q_{|X|}(i/d)| \rightarrow 0$ in probability, verifying A4.2 with $\varphi = Q_{|X|}$. Under infinite width regime, conditional on x the pre-activations $h_j^{(1)}$ are approximately Gaussian with variance $\|x\|_2^2/d$, yielding a deterministic limit law for $\psi_{n,j}$ and hence A4.3. Moreover, although $\theta_{ij}^{(1)}$ appears inside $\psi_{n,j}$ through $h_j^{(1)}$, the influence of a single entry $\theta_{ij}^{(1)}$ on $h_j^{(1)}$ vanishes asymptotically as $d \rightarrow \infty$, satisfying A4.4. Finally, taking τ_n as the empirical $(1 - \rho)$ -quantile of $\{(S_n)_{ij}\}_{i,j}$ yields A4.5 under mild continuity conditions. We provide full verifications in the Appendix B.

4.3. Graphon convergence of factorised masks

Theorem 4.7 (Bipartite graphon convergence). *Let $d \rightarrow \infty$ and $n \rightarrow \infty$. Consider the factorised saliency model*

$$(S_n)_{ij} = \varphi_{n,i} \cdot \psi_{n,j} \cdot |\xi_{n,ij}|, \quad (M_n)_{ij} = \mathbf{1}\{(S_n)_{ij} > \tau_n\},$$

where $i \in [d]$, $j \in [n]$, and let $W_n := W_{M_n}$ be the associated bipartite step-kernel on $[0, 1]^2$. Under assumptions 4.1 - 4.5, we define the deterministic limit graphon

Table 1. Factorisation of PaI methods. SNIP (Lee et al., 2019b) and GraSP (Wang et al., 2020) are different in global scaling.

Method	Input φ_i	Neuron ψ_j	Noise $ \xi_{ij} $
SNIP	$ x_i $	$ \theta_j^{(2)} \sigma'(h_j^{(1)}) $	$ \theta_{ij}^{(1)} $
GraSP	$ x_i $	$ \theta_j^{(2)} \sigma'(h_j^{(1)}) $	$ \theta_{ij}^{(1)} $
Synflow	1	$ \theta_j^{(2)} $	$ \theta_{ij}^{(1)} $
Magnitude	1	1	$ \theta_{ij}^{(1)} $
Random	1	1	Unif $[0, 1]$

$$\mathcal{W}(u, v) := \mathbb{P}(\varphi(u) \cdot Q_\psi(v) \cdot |\xi| > \tau), \quad (u, v) \in [0, 1]^2.$$

Then $\delta_{\square}^{\text{bip}}(W_n, \mathcal{W}) \xrightarrow{P} 0$.

Remark 4.8. The limit graphon \mathcal{W} can be viewed as a topological signature of a PaI method. Unstructured pruning rules, such as Random or Magnitude pruning, converge to constant graphons $\mathcal{W}(u, v) \equiv \rho$, reflecting uniform connectivity across inputs and neurons. In contrast, data-dependent PaI methods induce heterogeneous graphons, where the probability of retaining a connection varies smoothly with latent input features and neuron sensitivities. From this perspective, different PaI rules are distinguished not by individual edge realisations, but by the large-scale connectivity patterns they impose in the infinite-width limit.

Remark 4.9. This convergence result extends naturally to deep architectures. For an L -hidden-layer network with widths (n_0, n_1, \dots, n_L) , each layer $\ell \in \{1, \dots, L\}$ has mask $M^{(\ell)} \in \{0, 1\}^{n_{\ell-1} \times n_\ell}$ generated by factorised scores in Eqn. (3). The corresponding bipartite step-kernel is $W_{M_n^{(\ell)}}$ or $W_n^{(\ell)}$ for short. Here $\varphi_{n,i}^{(\ell)}$ is a forward feature determined by the activation of neuron i in layer $\ell - 1$, while $\psi_{n,j}^{(\ell)}$ is a backward feature determined by the normalised backpropagated sensitivity of neuron j in layer ℓ . Under standard infinite-width moment assumptions, these row and column features have deterministic empirical limits after sorting, and the edge noise is asymptotically decoupled from the neuron feature by a leave-one-out argument. Therefore each layer admits a deterministic limit graphon $W^{(\ell)}$. Applying Theorem 4.7 to each layer ℓ gives the graphon convergence for each layer $W_n^{(\ell)} \xrightarrow{P} W^{(\ell)}$. Please refer to Appendix D for details.

4.4. Empirical verification

We empirically illustrate Theorem 4.7 by comparing finite-width masks, produced by the original SNIP, to the predicted limit graphon. For visualisation, we generate masks $M_n \in \{0, 1\}^{n \times n}$ at fixed density $\rho = 0.2$ for $n \in \{200, 500, 1000, 2000, 4000\}$. Following Pham et al. (2025), we sort rows/columns by their latent factors (φ_i for rows and ψ_j for columns) and average the sorted masks over 100 seeds to obtain an empirical edge-probability matrix.

Besides, we compute the limit theoretical graphon, $\mathcal{W}(u, v)$, as a continuous function on the unit square $[0, 1]^2$. To




Figure 1. Visual convergence to the graphon limit. We compare averaged empirical masks (over 100 seeds) at increasing widths ($n = 200, 500, 1000, 2000, 4000$) against the analytically computed Theoretical Graphon. The density is fixed at $\rho = 0.2$.

make the limit object computable, we consider i.i.d. standard normal inputs and weights initialisation. Since inputs $x \sim \mathcal{N}(0, 1)$, the magnitude $|x|$ follows a standard Half-Normal (HN) distribution. We define $\varphi(u) = F_{\text{HN}}^{-1}(u) = \sqrt{2}\text{erf}^{-1}(u)$. The column factor is defined as $\psi = |\theta^{(2)}| |\sigma'(h)|$, where $\theta^{(2)} \sim \mathcal{N}(0, 1)$. We determine the quantile function $Q_\psi(v)$ for this product distribution via Monte Carlo simulation ($N = 10^6$ samples). The edge noise $|\xi_{ij}|$ is Half-Normal distributed with scale $\sigma_{\theta^{(1)}} = 1$ where $\theta^{(1)} \sim \mathcal{N}(0, 1)$. Given a global threshold τ (solved numerically to match density ρ), the theoretical edge probability is $\mathcal{W}_{(u,v)} = \mathbb{P}(\varphi(u)\psi(v)|\xi| > \tau) = 1 - \text{erf}(\frac{\tau}{\varphi(u)Q_\psi(v)\sqrt{2}})$.

Figure 1 shows that the averaged empirical masks become progressively smoother and rapidly approach the theoretical \mathcal{W} as n increases; by $n \geq 2000$ the empirical and theoretical patterns are visually indistinguishable. This supports the interpretation of PaI masks as finite-sample discretisations of a deterministic graphon limit.

5. Expressivity

This section studies the approximation power of single-hidden-layer networks whose sparse connectivity is given by PaI masks. Classical Universal Approximation Theorems (UAT) rely on dense connectivity (Cybenko, 1989; Barron, 1994). For sparse networks, proving expressivity is highly non-trivial because algorithmic pruning is a random, combinatorial process that can abruptly destroy necessary computational paths.

Because the ambient input dimension satisfies $d \rightarrow \infty$, claiming a sparse model can approximate every continuous function of all d coordinates is both too strong and misaligned with real-world learning problems. Instead, we adopt an intrinsic-dimension regime (Nakada & Imaizumi, 2020), where the target function depends only on a fixed (or slowly growing) subset of k informative input coordinates. We ask: *does the pruned network retain universal approximation capabilities on this active subspace?*

Another challenge is that PaI masks are random, method-dependent *combinatorial* objects, so it is non-trivial to make precise connectivity claims directly at the level of the binary matrix M_n . Our graphon limit theory provides a deterministic continuum proxy to answer this: when the step-kernels W_{M_n} converge to a limit graphon \mathcal{W} , large-scale connec-

tivity properties of M_n can be stated as simple analytic conditions on \mathcal{W} . In particular, we demonstrate that if the limit graphon \mathcal{W} concentrates sufficient edge probability on an active subspace, then as $n \rightarrow \infty$, with high probability the finite mask M_n contains a fully connected $k \times \tilde{n}$ sub-network (with $\tilde{n} \rightarrow \infty$), allowing classical UAT (Cybenko, 1989) to apply on that core.

Masked 1-hidden-layer networks and deterministic threshold masks. Fix $d, n \in \mathbb{N}$. Let $M \in \{0, 1\}^{d \times n}$ be a binary mask sampled from a given graphon \mathcal{W} of the factorised saliency model studied in Sect. 4. We consider the class of masked two-layer neural networks

$$\mathcal{F}_n(M) := \left\{ x \mapsto \sum_{j=1}^n a_j \sigma \left(b_j + \sum_{i=1}^d M_{ij} \theta_{ij} x_i \right) \right\}, \quad (7)$$

where $(a_j), (b_j) \in \mathbb{R}^n, (\theta_{ij}) \in \mathbb{R}^{d \times n}$. Any conventional normalisation (e.g. $1/\sqrt{n}$) can be absorbed into a_j 's and does not affect expressivity.

5.1. Graphon coordinates and active rectangle

By the same measure-preserving relabelling used in the graphon convergence Theorem 4.7, we work in coordinates $u_i := i/d$ and $v_j := j/n$ such that $\varphi_{n,i} \approx \varphi(u_i)$ and the sorted neuron features satisfy $\psi_{n,(j)} \approx Q_\psi(v_j)$. In these coordinates, the limiting edge-probability kernel takes the explicit form $\mathcal{W}(u, v) = \mathbb{P}(\varphi(u)Q_\psi(v)|\xi| > \tau)$, matching the limit object from Sect. 4. Equivalently, one may view u_i, v_j as i.i.d. uniform latents upto a random relabelling, since the cut distance is invariant under index permutations.

Class of target functions depending on k coordinates.

Fix $k \in \mathbb{N}$ and a compact set $K \subset \mathbb{R}^k$. For any set $I \subset [d]$ with $|I| = k$, we write $x_I \in \mathbb{R}^k$ for the subvector of $x \in \mathbb{R}^d$ restricted to coordinates in I . Given $f \in C(K)$, we define the lifted target $\tilde{f} : \mathbb{R}^d \rightarrow \mathbb{R}$ by $\tilde{f}(x) := f(x_I)$, where $x \in \mathbb{R}^d$, and the lifted compact domain $\tilde{K} := \{x \in \mathbb{R}^d : x_I \in K, \|x_{I^c}\|_\infty \leq B\}$ for some fixed $B < \infty$, where $I^c := [d] \setminus I$. Note that approximation on \tilde{K} is equivalent to approximation of f on K , as \tilde{f} ignores x_{I^c} .

Theorem 5.1 (UAT for dense 1-hidden-layer neural networks on \mathbb{R}^k , (Cybenko, 1989)). *Fix $k \in \mathbb{N}$ and a compact $K \subset \mathbb{R}^k$. Suppose σ is a universal activation. Then the class $\{u \mapsto \sum_{r=1}^n a_r \sigma(\theta_r^\top u + b_r) : n \in \mathbb{N}\}$ is dense in*

$C(K)$. In particular, for every $f \in C(K)$ and every $\varepsilon > 0$, there exist an integer \tilde{n}_ε and parameters $\{(a_r, \theta_r, b_r)\}_{r=1}^{\tilde{n}_\varepsilon}$ such that $\sup_{u \in K} \left| \sum_{r=1}^{\tilde{n}_\varepsilon} a_r \sigma(\theta_r^\top u + b_r) - f(u) \right| < \varepsilon$.

Assumption 5.2 (Lower-bounded core connectivity). There exist measurable sets $U_\star, V_\star \subset (0, 1)$ with Lebesgue measures $\beta := \lambda(U_\star) > 0$ and $\alpha := \lambda(V_\star) > 0$, and a constant $p_\star > 0$ such that

$$\inf_{u \in U_\star, v \in V_\star} \mathcal{W}(u, v) \geq p_\star \quad (8)$$

In the factorised model in Eqn. (3), a sufficient condition is that $\varphi(u) \geq \varphi_0 > 0$ on U_\star and $Q_\psi(v) \geq \psi_0 > 0$ on V_\star , in which case one may take $p_\star = \mathbb{P}(\varphi_0 \psi_0 |\xi| > \tau) > 0$.

Assumption 5.3 (Availability of k active inputs). Given a fixed k , with probability $1 - o(1)$ as $d \rightarrow \infty$, there exists an index set $I \subset [d]$ such that $|I| = k$ and $u_i \in U_\star, \forall i \in I$.

Assumption 5.4 (Core width growth). Let k be fixed. As $n \rightarrow \infty$, we have $(\alpha n) p_\star^k \rightarrow \infty$.

Remark 5.5. Assumption 5.2 is a *continuum* statement about the limiting connectivity pattern. It says that there is a positive-measure block of input latents U_\star and hidden-unit latents V_\star such that edges in that block survive with probability uniformly bounded below. For Random pruning (constant graphon $W \equiv \rho$), one may take $U_\star \times V_\star = [0, 1]^2$ and $p_\star = \rho$. For data-dependent PaI (e.g. SNIP-type), $U_\star \times V_\star$ can be interpreted as the region where the signal concentrates: U_\star corresponds to input coordinates with non-negligible forward saliency φ , and V_\star to hidden units with non-negligible backward/gradient feature Q_ψ .

5.2. Universal approximation on active coordinates

We show that the class $\mathcal{F}_n(M)$ is universal for approximating target functions depending on k active coordinates. The following lemma ensures the existence of a dense subnetwork of size $k \times \tilde{n}$ in the graphon-sparsified network.

Lemma 5.6 (Dense core in a graphon-sparsified mask). *Assume assumptions 5.2 -5.4 and let I be as in Assumption 5.3. Fix $\tilde{n} \in \mathbb{N}$. Then, with overwhelming probability, there exists $J \subset [n]$ with $|J| = \tilde{n}$ such that $M_{ij} = 1, \forall i \in I, \forall j \in J$.*

Theorem 5.7 (Universality on active coordinates). *Assume Theorem 5.1 and assumptions 5.2- 5.4. Fix $k \in \mathbb{N}$, a compact $K \subset \mathbb{R}^k$, a target $f \in C(K)$. For every $\varepsilon > 0$, there exist $B < \infty$ and an index set I as in assumption 5.3. Define \tilde{f} and \tilde{K} as above. Then, with high probability, there exists $F_n \in \mathcal{F}_n(M)$ such that $\sup_{x \in \tilde{K}} |F_n(x) - \tilde{f}(x)| < \varepsilon$*

Remark 5.8. Theorem 5.7 states that universality on any fixed k -coordinate subspace holds whenever the limit graphon \mathcal{W} contains an *active block* $U_\star \times V_\star$ with a uniform edge floor $\mathcal{W} \geq p_\star > 0$; equivalently, one needs $\alpha n (p_\star)^k \rightarrow \infty$ to ensure enough hidden units connect simultaneously to k active inputs. This result highlights the fundamental difference between pruning methodologies at

extreme sparsities. For Random pruning ($\mathcal{W} \equiv \rho$), this reduces to $n\rho^k \rightarrow \infty$, which can fail at high sparsity since edges are spread uniformly. In contrast, data-dependent PaI can concentrate the same sparsity budget on task-relevant coordinates, yielding a larger effective floor p_\star on a non-vanishing active ‘‘sub-graphon’’ $U_\star \times V_\star$ (with α, β bounded below), thereby better preserving dense cores and approximation power in the high-sparsity regime.

6. Generalisation error bound

Generalisation error bounds for overparameterised networks depend crucially on *architecture and training regime* (Novak et al., 2018; Arora et al., 2019a). For PaI, the architecture itself is random or data-dependent through masks M_n , making it unclear what the right limiting object is. The graphon limit framework makes this dependence precise: if $W_{M_n} \rightarrow \mathcal{W}$, then the associated Graphon-NTK admits a deterministic limit that depends only on \mathcal{W} and the data distribution (Pham et al., 2025). This lifts generalisation analysis from a random finite mask to a deterministic operator defined by \mathcal{W} . Here, we derive the generalisation error bound for overparameterised sparse neural networks in the Graphon-NTK regime, in parallel with the bounds for dense models in (Cao & Gu, 2019; Arora et al., 2019a). Then, we interpret this bound through the input-output path density of layer graphons, which shows how the connectivity pattern of pruning masks affects generalisation capacity.

6.1. From graphon limits of PaI to Graphon-NTK

Let \mathcal{D} be a data distribution and let $S = \{(x_i, y_i)\}_{i=1}^m$ be i.i.d. samples from \mathcal{D} . As discussed in Section 4, masks generated by PaI methods converge to specific graphons, which describes the infinite-width limit of the induced sparsity structure. Given such limiting graphon \mathcal{W} , the authors of (Pham et al., 2025) defines the associated *Graphon-NTK kernel* $\Theta_{\mathcal{W}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ as the deterministic infinite-width limit of the NTK for graphon-induced sparse networks. On the dataset S , we define the *Graphon-NTK Gram matrix* $(K_{\mathcal{W}})_{ij} := \Theta_{\mathcal{W}}(x_i, x_j) \in \mathbb{R}^{m \times m}$.

6.2. Graphon-NTK generalisation bound

We analyse the infinite-width *Graphon-NTK (lazy) regime*, in which training dynamics of the masked network on a fixed sample are governed by the limiting $K_{\mathcal{W}}$. This viewpoint parallels the classical NTK linearisation results (Lee et al., 2019a; Chizat et al., 2019). Under this regime, we can directly apply the generalisation bound of Cao & Gu (2019) with $K_{\mathcal{W}}$ in place of the finite-width NTK Gram matrix.

Corollary 6.1 (Infinite-width Graphon-NTK generalisation bound). *Fix $\delta \in (0, e^{-1})$ and assume $\lambda_{\min}(K_{\mathcal{W}}) \geq \lambda_0 > 0$. As widths $\rightarrow \infty$, with probability at least $1 - \delta$ over S , the expected classification error of an associated output*

predictor of the neural network \hat{f} (see Appendix F.3) satisfies

$$\mathbb{E}[\mathcal{L}_{0.1}^{\mathcal{D}}(\hat{f})] \leq CL \cdot \sqrt{\frac{y^\top K_{\mathcal{W}}^{-1} y}{m}} + C \sqrt{\frac{\log(1/\delta)}{m}}, \quad (9)$$

where $y = (y_1, \dots, y_m)^\top$, $C > 0$ is a universal constant, and L is the depth factor.

Remark 6.2. The complexity term $y^\top K_{\mathcal{W}}^{-1} y$ measures the RKHS norm of the label vector under kernel $\Theta_{\mathcal{W}}$. Smaller values indicate the target labels are easier to separate in the induced feature space, directly linking pruning topology to generalisation capability. We provide proof in Appendix F.

6.3. Path density proxy for the Graphon-NTK in 2-hidden-layer networks ($L = 2$)

Corollary 6.1 shows that the generalisation error bound depends on the complexity term $y^\top K_{\mathcal{W}}^{-1} y$. The effect of the pruning topology can be understood through the mass of surviving input-output paths.

A path-density proxy. To interpret how the pruning topology affects the Graphon-NTK complexity term in Corollary 6.1, we consider $L = 2$ hidden layers with graphon limits $\mathcal{W}^{(1)}$, $\mathcal{W}^{(2)}$, and $\mathcal{W}^{(3)}$ of layerwise pruning masks, where $u_0, u_1, u_2, u_3 \in [0, 1]$ index input coordinates, hidden neurons, and output positions. Assume the activation σ is twice differentiable, and then the absolute value Graphon-NTK Gram matrix satisfies that

$$|K_{\mathcal{W}}(x, x')| \lesssim \left(\sum_{\ell=0}^1 C_\ell \left| \int_0^1 \Sigma^{(\ell)}(u_\ell, x, x') du_\ell \right| \right) \left(\int_0^1 |\Sigma^{(0)}(u_0, x, x')| P_{\mathcal{W}}(u_0) du_0 \right) \quad (10)$$

where C_ℓ 's are constants, $\Sigma^\ell(u_\ell, x, x')$ is activation covariance, $P_{\mathcal{W}}(u_0) = \int_{[0,1]^3} \mathcal{W}^{(1)}(u_0, u_1) \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1$ is the path density from input u_0 to output determined by layer graphons, and $\Sigma^0(u_0, x, x')$ is the input data correlation (See Appendix G for more details).

Interpretation. The path density $P_{\mathcal{W}}(u_0)$ acts as a coordinate-dependent measure of how much architectural mass connects input coordinate u_0 to the output. In a dense or uniformly random-pruned network, the path density is approximately constant across u_0 . For example, if $\mathcal{W}^{(\ell)} \equiv \rho_\ell$, then $P_{\mathcal{W}}(u_0) \equiv \rho_1 \rho_2 \rho_3$ up to the corresponding boundary interpretation for scalar output. Thus Random pruning uniformly rescales the contribution of all input coordinates.

In contrast, data-dependent PaI methods can induce heterogeneous graphons. If their graphons allocate larger connectivity mass to coordinates with larger forward saliency and to hidden units with larger backward sensitivity, then $P_{\mathcal{W}}(u_0)$ becomes larger on informative input regions. In

this case, the Graphon-NTK can place more weight on input correlations that are relevant for the target labels. This can improve both: (i) *spectral conditioning*: better eigenvalue spread in $K_{\mathcal{W}}$, avoiding spectral collapse; and (ii) *label alignment*: smaller complexity $y^\top K_{\mathcal{W}}^{-1} y$ when paths align with informative features.

Caveat. Overly peaked (concentrated) graphons risk spectral collapse if connectivity becomes too sparse in critical regions. Effective designs should maintain a balance between the core active region and the lower bounded probability structure: high path density on signal coordinates with nonzero background connectivity for stability. This is empirically observed and verified in (Pham et al., 2023).

6.4. Empirical results

We validate these theoretical insights by comparing the complexity measure $y^\top K_{\mathcal{W}}^{-1} y$ of data-driven pruning (SNIP Graphon) versus Random Pruning (Constant Graphon). We use wide finite-width MLP proxies ($n = 4096$) trained on binary CIFAR-10. We vary the label noise (ratio of flipped labels) and the network densities $\rho \in \{0.2, 0.4, 0.7, 0.8\}$.

Figure 2 shows the complexity bound $y^\top K_{\mathcal{W}}^{-1} y$ increases monotonically with label noise across densities, confirming that the bound correctly captures task difficulty. The gap between topologies is most pronounced in the sparse regime. At $\rho = 0.2$ and $\rho = 0.4$, the constant-graphon baseline yields much larger values, consistent with a more ill-conditioned $K_{\mathcal{W}}$, whereas the SNIP graphon remains stable. At higher densities ($\rho \geq 0.7$), the two curves become close (and can slightly reverse), suggesting that once the network is sufficiently dense, the bound is less sensitive to the precise pruning topology.

To provide direct evidence linking our graphon path-density theory to practical performance, we evaluate SNIP and Random pruning on ResNet-20 trained on CIFAR-10 across varying sparsities, visualised in the rightmost plot in Figure 2. We trained all these subnetworks for 160 epochs using the SGD optimiser with a learning rate that started at 0.1 and decayed during training. The result shows that Random is competitive with SNIP at medium sparsity, but degrades much more sharply at extreme sparsity. This is qualitatively consistent with the sensitivity result above and also the graphon/path-density picture: homogeneous random pruning uniformly dilutes connectivity, which becomes harmful at high sparsity, while data-dependent pruning can preserve higher path mass on informative coordinates.

Takeaway. The graphon limit offers a convenient way to express how the pruning topology enters the generalisation bound through the complexity term. This offers a way to understand the behaviour of PaI methods. In particular, under low density (high sparsity), data-dependent PaI tends to exhibit a lower $y^\top K_{\mathcal{W}}^{-1} y$, consistent with more concen-




Figure 2. Sensitivity of Graphon-NTK Complexity vs. Label Noise and Empirical Pruning Performance. The first four panels plot the theoretical complexity $y^\top K_W^{-1} y$ (y -axis) against the ratio of randomised labels (x -axis) for a specific density ρ . The rightmost panel shows the empirical test accuracy against sparsity for SNIP and Random pruning on CIFAR-10 with (Wide)ResNet-20.

trated (heterogeneous) path allocation. This explains why SNIP/GraSP perform better than Random pruning. In contrast, under mild pruning, the constant-graphon (Random pruning) yields a comparable or even smaller complexity term, matching its often competitive performance.

7. Conclusion

In this work, we provide the first rigorous derivation of the infinite-width limits for PaI methods. By introducing the Factorised Saliency Model, we prove that discrete masks generated by popular PaI methods converge to limit bipartite graphons. This limit captures the asymptotic sparsity geometry and yields a principled way to compare different pruning rules through their induced graphons. In particular, Random pruning converges to homogeneous graphons, and data-driven methods induce structured heterogeneity. Leveraging this graphon limit perspective, we establish two fundamental guarantees for sparse learning. First, we derive a Universal Approximation Theorem for sparse networks restricted to “active” coordinate subspaces, demonstrating that sparse networks retain the expressivity of dense models on the manifold of relevant features. Second, through the extension of the Graphon-NTK, we link topology to generalisation via the concept of path density. Our analysis confirms that heterogeneous graphons improve kernel alignment by concentrating path density on informative signals, whereas homogeneous ones dilute path density uniformly. Future work could leverage these topological insights to design principled pruning criteria that focus on maximising path density on important features. Additionally, investigating the finite-width corrections to this asymptotic theory would help further bridge the gap to practical implementation.

Limitations and future works. While our graphon-limit framework provides a rigorous mathematical foundation for analysing sparse topologies, we acknowledge several theoretical limitations that present exciting future research.

Firstly, our graphon convergence theorem currently assumes a joint infinite-width limit where both the hidden width and input dimension approach infinity. If the input (or output) dimension is kept fixed, as is standard in classical NTK analyses, the mask would no longer converge to a standard

continuous graphon. Instead, the natural limit object becomes a mixed discrete-continuous kernel. Extending our framework to this regime is a crucial next step for aligning the theory closer to standard architectures.

Secondly, to mathematically verify the convergence of the input saliency profile (Assumption 4.2), our single-layer analysis utilises bounded i.i.d. inputs. While this is a strong assumption for structured raw data (e.g., images), it is intentionally designed to be transferable to hidden layers. As outlined in our layer-wise blueprint for deep networks (Appendix D), hidden activations naturally converge to deterministic profiles under standard infinite-width limits, thereby organically satisfying this assumption at deeper layers. Nevertheless, establishing a fully rigorous deep PaI theorem remains an open challenge. Because pruning decisions at early layers perturb the forward activations and backward gradients of subsequent layers, future work must establish tighter, non-asymptotic control over these cross-layer dependencies.

Thirdly, our results all operate in the strict infinite-width limit. In practice, algorithmically pruned networks operate at finite widths where combinatorial variance still impacts the topology. Therefore, a highly impactful future direction is the derivation of finite-width convergence rates. Developing non-asymptotic bounds that quantify how fast empirical masks approach the graphon, and precisely how approximation power and NTK conditioning degrade at finite widths under sparsity, will further bridge the gap between asymptotic theory and practical implementation.

Impact Statement

This work establishes a rigorous continuous-limit framework for Pruning-at-Initialisation (PaI), characterising the asymptotic limits of sparse network topologies via continuous function kernels. We prove that popular PaI methods converge to deterministic graphons, providing the mathematical tools guarantee the expressivity and generalisation of sparse networks. These insights pave the way for designing principled, highly efficient sparse training algorithms, directly contributing to the effort to reduce the computational and environmental costs of modern deep learning.

References

- Arora, S., Du, S., Hu, W., Li, Z., and Wang, R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International conference on machine learning*, pp. 322–332. PMLR, 2019a.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R., and Wang, R. On exact computation with an infinitely wide neural net. *Advances in neural information processing systems*, 32, 2019b.
- Barron, A. R. Approximation and estimation bounds for artificial neural networks. *Machine learning*, 14(1):115–133, 1994.
- Bolcskei, H., Grohs, P., Kutyniok, G., and Petersen, P. Optimal approximation with sparsely connected deep neural networks. *SIAM Journal on Mathematics of Data Science*, 1(1):8–45, 2019.
- Borgs, C., Chayes, J. T., Lovász, L., Sós, V. T., and Veszteg, K. Convergent sequences of dense graphs i: Subgraph frequencies, metric properties and testing. *Advances in Mathematics*, 219(6):1801–1851, 2008.
- Burkholz, R. Convolutional and residual networks provably contain lottery tickets. In *International Conference on Machine Learning*, pp. 2414–2433. PMLR, 2022a.
- Burkholz, R. Most activation functions can win the lottery without excessive depth. *Advances in Neural Information Processing Systems*, 35:18707–18720, 2022b.
- Cao, Y. and Gu, Q. Generalization bounds of stochastic gradient descent for wide and deep neural networks. *Advances in neural information processing systems*, 32, 2019.
- Chizat, L., Oyallon, E., and Bach, F. On lazy training in differentiable programming. *Advances in neural information processing systems*, 32, 2019.
- Cybenko, G. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Eldan, R. and Shamir, O. The power of depth for feedforward neural networks. In *Conference on learning theory*, pp. 907–940. PMLR, 2016.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pp. 3259–3269. PMLR, 2020.
- Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Pruning neural networks at initialization: Why are we missing the mark? In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ig-VyQc-MLK>.
- Gebhart, T., Saxena, U., and Schrater, P. A unified paths perspective for pruning at initialization. *ArXiv*, abs/2101.10552, 2021.
- Han, S., Pool, J., Tran, J., and Dally, W. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015.
- Hoang, D. N., Liu, S., Marculescu, R., and Wang, Z. REVISITING PRUNING AT INITIALIZATION THROUGH THE LENS OF RAMANUJAN GRAPH. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=uVcDssQff_.
- Hornik, K. Approximation capabilities of multi-layer feedforward networks. *Neural Networks*, 4(2):251–257, 1991. ISSN 0893-6080. doi: [https://doi.org/10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T). URL <https://www.sciencedirect.com/science/article/pii/089360809190009T>.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- Keriven, N., Bietti, A., and Vaiter, S. Convergence and stability of graph convolutional networks on large random graphs. *Advances in Neural Information Processing Systems*, 33:21512–21523, 2020.
- Krishnagopal, S. and Ruiz, L. Graph neural tangent kernel: Convergence on large graphs. In *International Conference on Machine Learning*, pp. 17827–17841. PMLR, 2023.
- Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., and Pennington, J. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019a.
- Lee, N., Ajanthan, T., and Torr, P. Snip: Single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*, 2019b. URL <https://openreview.net/forum?id=B1VZqjAcYX>.
- Liu, S., Chen, T., Chen, X., Shen, L., Mocanu, D. C., Wang, Z., and Pechenizkiy, M. The unreasonable effectiveness of random pruning: Return of the most naive baseline for

- sparse training. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=VBZJ_3tz-t.
- Liu, S., Cheng, Y., Zha, F., Guo, W., Sun, L., Bing, Z., and Yang, C. Learning effective pruning at initialization from iterative pruning. *arXiv preprint arXiv:2408.14757*, 2024.
- Lovász, L. and Szegedy, B. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96(6):933–957, 2006.
- Malach, E., Yehudai, G., Shalev-Schwartz, S., and Shamir, O. Proving the lottery ticket hypothesis: Pruning is all you need. In *International Conference on Machine Learning*, pp. 6682–6691. PMLR, 2020.
- Nakada, R. and Imaizumi, M. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *Journal of Machine Learning Research*, 21(174):1–38, 2020.
- Novak, R., Bahri, Y., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Sensitivity and generalization in neural networks: an empirical study. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HJC2SzZCW>.
- Orseau, L., Hutter, M., and Rivasplata, O. Logarithmic pruning is all you need. *Advances in Neural Information Processing Systems*, 33:2925–2934, 2020.
- Patil, S. M. and Dovrolis, C. Phew: Constructing sparse networks that learn fast and generalize well without training data. In *International Conference on Machine Learning*, pp. 8432–8442. PMLR, 2021.
- Pham, H., Liu, S., Xiang, L., Le, D., Wen, H., Tran-Thanh, L., et al. Towards data-agnostic pruning at initialization: what makes a good sparse mask? *Advances in Neural Information Processing Systems*, 36:80044–80065, 2023.
- Pham, H., Ta, T.-A., Jacobs, T., Burkholz, R., and Tran-Thanh, L. The graphon limit hypothesis: Understanding neural network pruning via infinite width analysis. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=EEZLBhyer1>.
- Ruiz, L., Chamon, L., and Ribeiro, A. Graphon neural networks and the transferability of graph neural networks. *Advances in Neural Information Processing Systems*, 33: 1702–1712, 2020.
- Su, J., Chen, Y., Cai, T., Wu, T., Gao, R., Wang, L., and Lee, J. D. Sanity-checking pruning methods: Random tickets can win the jackpot. *Advances in Neural Information Processing Systems*, 33:20390–20401, 2020.
- Tanaka, H., Kunin, D., Yamins, D. L., and Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in Neural Information Processing Systems*, 33:6377–6389, 2020.
- Tropp, J. A. et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- Wang, C., Zhang, G., and Grosse, R. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkgsACVKPH>.
- Xiang, L., Nguyen-Tri, Q., Nguyen, L.-C., Pham, H., Than, K., Tran-Thanh, L., and Wen, H. DPai: Differentiable pruning at initialization with node-path balance principle. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=hvLBTpiDt3>.
- Yang, H. and Wang, Z. On the neural tangent kernel analysis of randomly pruned neural networks. In Ruiz, F., Dy, J., and van de Meent, J.-W. (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 1513–1553. PMLR, 25–27 Apr 2023. URL <https://proceedings.mlr.press/v206/yang23b.html>.
- Zhou, H., Lan, J., Liu, R., and Yosinski, J. Deconstructing lottery tickets: Zeros, signs, and the supermask. *Advances in neural information processing systems*, 32, 2019.

Organisation of the Appendix. This appendix provides detailed mathematical proofs and additional supporting experiments for the results in the main text.

- Appendix A provides the proof of Theorem 4.7 (Graphon Convergence) using the Factorised Saliency Model.
- Appendices B and C provide formal verification of the fact that SNIP and GraSP satisfy the regularity assumptions required for application of Theorem 4.7.
- Appendix D provides details on extension of the Graphon Convergence Theorem to deep neural networks.
- Appendix E provides the proof of Theorem 5.7 (Universal Approximation on Active Subspaces).
- Appendices F and G provide the proofs for the Corollary 6.1 (Generalisation error bound) and the path density decomposition.

A. Graphon Convergence for Factorised Saliency Model

In this appendix, we prove the graphon convergence in bipartite cut-metric for the factorised saliency model

$$S_{n,ij} = \varphi_{n,i} \psi_{n,j} |\xi_{n,ij}|, \quad M_{n,ij} = \mathbf{1}\{S_{n,ij} > \tau_n\}, \quad i \in [d], j \in [n],$$

in the joint limit $d = d(n) \rightarrow \infty$ and $n \rightarrow \infty$. Let W_A denote the bipartite step kernel associated with a matrix $A \in \mathbb{R}^{d \times n}$, and $\delta_{\square}^{\text{bip}}$ denote the bipartite cut distance Eqn. (2). We first restate assumptions and main results, then provide detailed explanations and proofs.

A.1. Assumptions

We will take the joint limit as $d \rightarrow \infty$ and $n \rightarrow \infty$. We state the assumptions at the necessary level of generality needed for our cut-metric convergence argument. We add remarks to indicate typical sufficient conditions.

Assumption A.1 (Asymptotic width scaling).

$$\frac{\log(d+n)}{\min\{d, n\}} \rightarrow 0.$$

Assumption A.1 ensures that uniform concentration bounds over cut sets dominate logarithm of the cardinality terms arising from discretisation and union bounds. It is satisfied, for instance, when d and n grow at any polynomial rate.

Assumption A.2 (Deterministic input features). There exists a bounded continuous function $\varphi : [0, 1] \rightarrow [0, \infty)$ such that, after a relabeling of input coordinates it holds that

$$\max_{1 \leq i \leq d} |\varphi_{n,i} - \varphi(i/d)| \xrightarrow{\mathbb{P}} 0.$$

This assumption identifies a continuum limit for the discrete row features and guarantees that the row part of the edge-probability grid is well-approximated by $\varphi(u)$ uniformly in i . Continuity of φ controls discretization error when passing from i/d to u .

Assumption A.3 (Neuron features empirical CDF convergence). Let $\widehat{F}_{\psi,n}(t) := \frac{1}{n} \sum_{j=1}^n \mathbf{1}\{\psi_{n,j} \leq t\}$. There exists a deterministic CDF F_{ψ} supported on $[0, \infty)$ such that

$$\sup_{t \in \mathbb{R}} |\widehat{F}_{\psi,n}(t) - F_{\psi}(t)| \xrightarrow{\mathbb{P}} 0.$$

Let $Q_{\psi}(v) := \inf\{y \geq 0 : F_{\psi}(y) \geq v\}$ be the generalised quantile. After sorting columns by $\psi_{n,j}$, the kernel becomes a deterministic function of the order statistics $\psi_{n,j}$. This assumption ensures that the sorted column features converge via order statistics to the quantile curve $v \mapsto Q_{\psi}(v)$, which defines the limiting column coordinate.

Assumption A.4 (Conditional independence of edge noise). Let

$$\mathcal{G}_n := \sigma(\{\varphi_{n,i}\}_{i \leq d}, \{\psi_{n,j}\}_{j \leq n}, \tau_n).$$

Conditioned on \mathcal{G}_n , the noises $\{\xi_{n,ij}\}_{i \leq d, j \leq n}$ are independent with common law P_{ξ} , and are independent of \mathcal{G}_n . Moreover, the CDF of $|\xi|$ on $[0, \infty)$ is continuous.

In this assumption, independence of $\xi_{n,ij}$ yields conditional independence of mask entries given $(\varphi_n, \psi_n, \tau_n)$, enabling application of the matrix Bernstein concentration for the random mask around its conditional mean. Continuity of the CDF of $|\xi|$ ensures that the link function $z \mapsto \mathbb{P}(z|\xi| > \tau)$ is continuous, which is required to pass to the limit in the mean kernel.

Assumption A.5 (Deterministic threshold). There exists $\tau \in (0, \infty)$ such that

$$\tau_n \xrightarrow{\mathbb{P}} \tau \quad \text{and} \quad \tau_n \text{ is } \mathcal{G}_n\text{-measurable.}$$

This assumption allows the threshold to control sparsity. Convergence of τ_n ensures that the edge-probability link function stabilizes asymptotically. Measurability with respect to \mathcal{G}_n guarantees that conditioning on \mathcal{G}_n indeed implies edges independent via Assumption A.4.

A.2. Proof of Theorem 4.7

Theorem A.6 (Bipartite graphon convergence, Theorem 4.7). *Let $d \rightarrow \infty$ and $n \rightarrow \infty$. Consider the model*

$$S_{n,ij} = \varphi_{n,i} \cdot \psi_{n,j} \cdot |\xi_{n,ij}|, \quad M_{n,ij} = \mathbf{1}\{S_{n,ij} > \tau_n\}, \quad i \in [d], j \in [n], \quad (11)$$

and let $W_n := W_{M_n}$ denote the step kernel associated with M_n . Assume assumptions A.1–A.5. Define the limiting kernel $\mathcal{W} : [0, 1]^2 \rightarrow [0, 1]$ by

$$\mathcal{W}(u, v) := \mathbb{P}(\varphi(u) \cdot Q_\psi(v) \cdot |\xi| > \tau), \quad (12)$$

where $Q_\psi := F_\psi^{-1}(v) := \inf\{y : F_\psi(y) \geq v\}$ is the generalized quantile function. Then

$$\delta_{\square}^{\text{bip}}(W_n, \mathcal{W}) \xrightarrow{\mathbb{P}} 0. \quad (13)$$

Remark A.7. The limiting kernel \mathcal{W} is monotonic in v due to the use of the quantile function $Q_\psi(v)$.

High-level summary of the proof. We first sketch the main ideas of the proof. Detailed technical steps will be given later.

Let π_n be a sorting permutation of the columns by $\psi_{n,j}$. We write W_n^π for the corresponding column-relabeling which does not change $\delta_{\square}^{\text{bip}}$. Let \bar{W}_n^π be the conditional-mean kernel, whose matrix entries are $\bar{M}_{n,ij}^\pi = \mathbb{P}(\varphi_{n,i} \psi_{n,\pi_n(j)} |\xi| > \tau_n \mid \varphi_n, \psi_n, \tau_n)$. Then

$$\|W_n^\pi - \mathcal{W}\|_{\square} \leq \|W_n^\pi - \bar{W}_n^\pi\|_{\square} + \|\bar{W}_n^\pi - \mathcal{W}\|_{\square}. \quad (14)$$

(i) *Step 1: Concentration bound* (proved in A.3.3) Conditioned on $(\varphi_n, \psi_n, \tau_n)$, the centered entries $M_{n,ij}^\pi - \bar{M}_{n,ij}^\pi$ are independent, mean-zero, and bounded, so an application of the matrix Bernstein bound in combination with the estimate $\|\cdot\|_{\square} \lesssim (dn)^{-1/2} \|\cdot\|_{\text{op}}$ yields $\|W_n^\pi - \bar{W}_n^\pi\|_{\square} \xrightarrow{\mathbb{P}} 0$ under Assumption 4.1.

(ii) *Step 2: Mean-kernel limit* (proved in A.3.4) Assumptions 4.2–4.5 imply that $\bar{W}_n^\pi(u, v) \rightarrow \mathcal{W}(u, v)$. Using the boundedness of $(0 \leq \bar{W}_n^\pi \leq 1)$, it follows that $\|\bar{W}_n^\pi - \mathcal{W}\|_1 \xrightarrow{\mathbb{P}} 0$, hence $\|\bar{W}_n^\pi - \mathcal{W}\|_{\square} \xrightarrow{\mathbb{P}} 0$. Combining the two terms and using the invariance under relabelling of the cut metric gives the claim.

A.3. Technical Details for the Proof of Theorem 4.7

A.3.1. EQUIVALENCE OF CUT NORMS

Lemma A.8 (Step kernel cut norm equals matrix cut norm). *For $A \in \mathbb{R}^{d \times n}$, define the matrix cut norm*

$$\|A\|_{\square, \text{mat}} := \frac{1}{dn} \max_{S \subseteq [d], T \subseteq [n]} \left| \sum_{i \in S} \sum_{j \in T} A_{ij} \right|. \quad (15)$$

Then $\|W_A\|_{\square} = \|A\|_{\square, \text{mat}}$.

Proof. Any choice of subsets $S \subseteq [d]$ and $T \subseteq [n]$ corresponds to measurable sets $S^* = \bigcup_{i \in S} I_i$ and $T^* = \bigcup_{j \in T} J_j$. That yields $\|W_A\|_{\square} \geq \|A\|_{\square, \text{mat}}$.

Conversely, let $S, T \subset [0, 1]$ be measurable sets. Let $s_i := d\mu(S \cap I_i) \in [0, 1]$ and $t_j := n\mu(T \cap J_j) \in [0, 1]$, where μ is the Lebesgue measure. Since W_A is constant on each set $I_i \times J_j$

$$\int_{S \times T} W_A = \frac{1}{dn} \sum_{i=1}^d \sum_{j=1}^n A_{ij} s_i t_j. \quad (16)$$

For fixed $t = (t_j)$, the expression is linear in each $s_i \in [0, 1]$, so it is maximized at an extreme point $s_i \in \{0, 1\}$. The maximum over $[0, 1]^{d+n}$ of a multilinear form is achieved at a vertex. Similarly for t_j . Hence the supremum over measurable S, T equals the maximum over $S_0 \subset [d], T_0 \subset [n]$, i.e. $\|W_A\|_{\square} \leq \|A\|_{\square, \text{mat}}$. \square

By Lemma A.8, we can work directly with matrix cut norms.

A.3.2. COLUMN SORTING AND INVARIANCE

Let π_n be a permutation such that $\psi_{n, \pi_n(1)} \leq \psi_{n, \pi_n(2)} \leq \dots \leq \psi_{n, \pi_n(n)}$. We define the column-permuted mask matrix and its step kernel as

$$M_{n, ij}^{\pi} := M_{n, i\pi_n(j)}, \quad W_n^{\pi} := W_{M_n^{\pi}}. \quad (17)$$

Lemma A.9 (Invariance under row and column permutation). *We have*

$$\delta_{\square}^{\text{bip}}(W_n, \mathcal{W}) = \delta_{\square}^{\text{bip}}(W_n^{\sigma, \pi}, \mathcal{W}) \leq \|W_n^{\sigma, \pi} - \mathcal{W}\|_{\square}. \quad (18)$$

Proof. A row and column permutations correspond to a measure-preserving relabelling of the u - and v -axes that maps each interval I_i to $I_{\sigma_n(i)}$ and J_j to $J_{\pi_n(j)}$, respectively. Since the bipartite cut distance permits independent relabelings of rows and columns, it is invariant under this transformation. The final inequality follows by taking the trivial identity transformation in the infimum defining $\delta_{\square}^{\text{bip}}$. \square

Lemma A.9 reduces the problem to showing $\|W_n^{\pi} - \mathcal{W}\|_{\square} \rightarrow 0$. To control this distance, we introduce an intermediate proxy, which is the conditional-mean kernel, denoted by \bar{W}_n^{π} . From the assumption A.4, conditioned on \mathcal{G}_n , the entries $M_{n, ij}^{\pi}$ are conditionally independent Bernoulli random variables with conditional means

$$\bar{M}_{n, ij}^{\pi} := \mathbb{E}[M_{n, ij}^{\pi} \mid \mathcal{G}_n] = \mathbb{P}(\varphi_{n, i} \cdot \psi_{n, \pi_n(j)} \mid \xi > \tau_n). \quad (19)$$

Let $\bar{W}_n^{\pi} := W_{\bar{M}_n^{\pi}}$ be the step kernel induced by \bar{M}_n^{π} . By the triangle inequality, we decompose the error into a stochastic concentration term and a deterministic approximation term:

$$\|W_n^{\pi} - \mathcal{W}\|_{\square} \leq \|W_n^{\pi} - \bar{W}_n^{\pi}\|_{\square} + \|\bar{W}_n^{\pi} - \mathcal{W}\|_{\square} \quad (20)$$

We analyse these two terms separately in Appendices A.3.3, A.3.4, respectively.

A.3.3. CONCENTRATION AROUND CONDITIONAL MEAN

The term $\|W_n^{\pi} - \bar{W}_n^{\pi}\|_{\square}$ captures the deviation of the binary edges from their conditional probabilities due to the noise ξ . We define the difference matrix

$$X_n := M_n^{\pi} - \bar{M}_n^{\pi} \in \mathbb{R}^{d \times n}. \quad (21)$$

Conditioned on \mathcal{G}_n , we have

- The entries $\{X_{n, ij}\}_{i \leq d, j \leq n}$ are independent, since they depend on i.i.d. $\{\xi_{n, ij}\}$, and π_n is a deterministic function of ψ_n ,
- $\mathbb{E}[X_{n, ij} \mid \mathcal{G}_n] = 0$,
- $|X_{n, ij}| \leq 1$.

Our goal is to show that $\|W_n^\pi - \bar{W}_n^\pi\|_\square = \|X_n\|_{\square, \text{mat}} \xrightarrow{\mathbb{P}} 0$.

Lemma A.10 (Cut norm controlled by operator norm). *For any matrix $B \in \mathbb{R}^{d \times n}$,*

$$\|B\|_{\square, \text{mat}} \leq \frac{1}{\sqrt{dn}} \|B\|_{\text{op}}. \quad (22)$$

Proof. By definition,

$$\|B\|_{\square, \text{mat}} = \frac{1}{dn} \max_{S \subseteq [d], T \subseteq [n]} \left| \sum_{i \in S} \sum_{j \in T} B_{ij} \right|. \quad (23)$$

For any S, T , let $s \in \{0, 1\}^d$ and $t \in \{0, 1\}^n$ be their indicator vectors. Then

$$\sum_{i \in S} \sum_{j \in T} B_{ij} = s^\top B t. \quad (24)$$

By the Cauchy-Schwarz inequality and the definition of operator norm,

$$|s^\top B t| \leq \|B\|_{\text{op}} \|s\|_2 \|t\|_2. \quad (25)$$

Since $\|s\|_2 = \sqrt{|S|} \leq \sqrt{d}$ and $\|t\|_2 = \sqrt{|T|} \leq \sqrt{n}$,

$$|s^\top B t| \leq \|B\|_{\text{op}} \sqrt{dn}. \quad (26)$$

Dividing by dn and maximising over S, T yields the result. \square

By Lemma A.10 and Lemma A.8, it suffices to prove $\|X_n\|_{\text{op}} = o(\sqrt{dn})$.

Lemma A.11 (Matrix Bernstein implies cut-norm concentration). *Under the assumption A.1, we have*

$$\|W_n^\pi - \bar{W}_n^\pi\|_\square \xrightarrow{\mathbb{P}} 0. \quad (27)$$

Proof. By Lemma A.8, we have $\|W_n^\pi - \bar{W}_n^\pi\|_\square = \|X_n\|_{\square, \text{mat}}$. By Lemma A.10, we have $\|X_n\|_{\square, \text{mat}} \leq \frac{1}{\sqrt{dn}} \|X_n\|_{\text{op}}$. We apply the rectangular matrix Bernstein inequality (Tropp et al., 2015) to the sum of independent mean-zero random matrices representing the entries of $X_n \in \mathbb{R}^{d \times n}$, conditioned on ψ_n : for all $t \geq 0$,

$$\mathbb{P}(\|X\|_{\text{op}} \geq t \mid \mathcal{G}_n) = \mathbb{P}\left(\left\| \sum_{i,j} A_{ij} \right\|_{\text{op}} \geq t \mid \mathcal{G}_n\right) \leq (d+n) \exp\left(-\frac{t^2}{2(V+Lt/3)}\right). \quad (28)$$

where A_{ij} is a zero matrix except the entry (i, j) where it's value is X_{ij} . Therefore, each matrix A_{ij} is zero mean and $|A_{ij}| \leq 1$. The bounding term is $L = \|A_{ij}\|_{\text{op}} = 1$. In addition, $V = \max\{\|\sum_{i,j} \mathbb{E}[A_{ij} A_{ij}^\top | \mathcal{G}_n]\|_{\text{op}}, \|\sum_{i,j} \mathbb{E}[A_{ij}^\top A_{ij} | \mathcal{G}_n]\|_{\text{op}}\}$ is the norm of the matrix variance of the sum. Here, each entry of X_n is a centered Bernoulli, so $\mathbb{E}[X_{ij}^2 | \mathcal{G}_n] \leq 1/4$. Thus, $V \leq \frac{1}{4} \max(d, n)$, and we obtain the unconditional bound

$$\mathbb{P}(\|X_n\|_{\text{op}} \geq t) \leq (d+n) \exp\left(-\frac{t^2}{2\left(\frac{1}{4} \max(d, n) + t/3\right)}\right). \quad (29)$$

We now choose $t := C\left(\sqrt{d} + \sqrt{n}\right) \sqrt{\log(d+n)}$ with $C > 0$ sufficiently large. Note that t^2 is of order $(d+n) \log(d+n)$, which dominates $\max(d, n)$ for large n . More precisely, for large n ,

$$\frac{t^2}{2\left(\frac{1}{4} \max(d, n) + t/3\right)} \geq c \log(d+n) \quad (30)$$

for some constant $c > 0$ that can be made arbitrarily large by increasing C . Choosing C so that $c > 2$, we obtain $\mathbb{P}(\|X_n\|_{\text{op}} \geq t) \rightarrow 0$.

Therefore, by Lemma A.10,

$$\|X_n\|_{\square, \text{mat}} \leq \frac{1}{\sqrt{dn}} \|X_n\|_{\text{op}} \leq \frac{t}{\sqrt{dn}} = C \left(\sqrt{\frac{\log(d+n)}{n}} + \sqrt{\frac{\log(d+n)}{d}} \right) \quad (31)$$

which converges to 0 by Assumption A.1. By Lemma A.8, we have $\|W_n^\pi - \bar{W}_n^\pi\|_{\square} = \|X_n\|_{\square, \text{mat}} \xrightarrow{\mathbb{P}} 0$. \square

A.3.4. MEAN KERNEL CONVERGENCE

The term $\|\bar{W}_n^\pi - \mathcal{W}\|_{\square}$ measures the discretisation error between the finite grid of probabilities and the continuous limit graphon. In this step, we will bound the cut norm by the stronger L^1 norm: $\|\bar{W}_n^\pi - \mathcal{W}\|_{\square} \leq \|\bar{W}_n^\pi - \mathcal{W}\|_1$.

Define the functions

$$g(\tau, z) := \mathbb{P}(z|\xi| > \tau), \quad z \geq 0, \tau > 0, \quad (32)$$

with $g(\tau, 0) = 0$. Then

$$\bar{M}_{n,ij}^\pi = g(\tau_n, \varphi_{n,i} \psi_{n,\pi_n(j)}), \quad \mathcal{W}(u, v) = g(\tau, \varphi(u) Q_\psi(v)). \quad (33)$$

Here, we call this functions $g()$ a link function that capture the fact that edge probability depends on the product of forward and backward factors.

Lemma A.12 (Order statistics from empirical CDF convergence). *Assume Assumption A.3. Let $v \in (0, 1)$ be a continuity point of Q_ψ and $j_n := \lceil nv \rceil$. Then*

$$\psi_{n,(j_n)} \xrightarrow{\mathbb{P}} Q_\psi(v). \quad (34)$$

Proof. Let $q := Q_\psi(v)$ and fix $\eta > 0$. Continuity of Q_ψ at v implies $F_\psi(q - \eta) < v < F_\psi(q + \eta)$. Set

$$\Delta_\eta := \min\{v - F_\psi(q - \eta), F_\psi(q + \eta) - v\} > 0. \quad (35)$$

On the event $\{\sup_t |\hat{F}_{\psi,n}(t) - F_\psi(t)| \leq \Delta_\eta\}$, we have $\hat{F}_{\psi,n}(q - \eta) < v < \hat{F}_{\psi,n}(q + \eta)$. By the generalized inverse identity $\hat{F}_{\psi,n}^{-1}(v) = \psi_{n,\lceil nv \rceil}$, it holds that $q - \eta \leq \psi_{n,(j_n)} \leq q + \eta$. Therefore

$$\mathbb{P}(|\psi_{n,(j_n)} - q| > \eta) \leq \mathbb{P}\left(\sup_t |\hat{F}_{\psi,n}(t) - F_\psi(t)| > \Delta_\eta\right) \rightarrow 0, \quad (36)$$

which shows $\psi_{n,j_n} \xrightarrow{\mathbb{P}} q$. \square

Lemma A.13 (L^1 convergence of the mean kernel). *It holds that*

$$\|\bar{W}_n^\pi - \mathcal{W}\|_1 \xrightarrow{\mathbb{P}} 0, \quad \text{and} \quad \|\bar{W}_n^\pi - \mathcal{W}\|_{\square} \xrightarrow{\mathbb{P}} 0. \quad (37)$$

Proof. Let $D_n(u, v) := |\bar{W}_n^\pi(u, v) - \mathcal{W}(u, v)| \in [0, 1]$. Fix (u, v) such that v is a continuity point of Q_ψ . Let $i_n := \lceil du \rceil$ and $j_n := \lceil nv \rceil$. By Assumption A.2 and continuity of φ , we have $\varphi_{n,i_n} \xrightarrow{\mathbb{P}} \varphi(u)$. By Lemma A.12, $\psi_{n,(j_n)} \xrightarrow{\mathbb{P}} Q_\psi(v)$. By Assumption A.5, $\tau_n \xrightarrow{\mathbb{P}} \tau$. Since g is continuous and the limit is deterministic, the continuous mapping theorem yields

$$\bar{W}_n^\pi(u, v) = g(\tau_n, \varphi_{n,i_n} \psi_{n,(j_n)}) \xrightarrow{\mathbb{P}} g(\tau, \varphi(u) Q_\psi(v)) = \mathcal{W}(u, v), \quad (38)$$

so that $D_n(u, v) \xrightarrow{\mathbb{P}} 0$ for all u and all continuity points v of Q_ψ . As Q_ψ has at most countably many discontinuities, this holds for Lebesgue-a.e. (u, v) . Because $0 \leq D_n(u, v) \leq 1$, for any $\varepsilon > 0$, $\mathbb{E}[D_n(u, v)] \leq \varepsilon + \mathbb{P}(D_n(u, v) > \varepsilon)$, this boundedness implies $\mathbb{E}[D_n(u, v)] \rightarrow 0$ for Lebesgue-a.e. (u, v) when $D_n(u, v) \xrightarrow{\mathbb{P}} 0$. By dominated convergence theorem,

$$\mathbb{E}\|\bar{W}_n^\pi - \mathcal{W}\|_1 = \int_{[0,1]^2} \mathbb{E}[D_n(u, v)] du dv \rightarrow 0. \quad (39)$$

The Markov inequality yields $\|\bar{W}_n^\pi - \mathcal{W}\|_1 \xrightarrow{\mathbb{P}} 0$. Finally, $\|\cdot\|_{\square} \leq \|\cdot\|_1$ gives the cut-norm claim. \square

A.3.5. CONCLUSION

Combining the above results and the triangle inequality, we have

$$\|W_n^\pi - \mathcal{W}\|_{\square} \leq \|W_n^\pi - \bar{W}_n^\pi\|_{\square} + \|\bar{W}_n^\pi - \mathcal{W}\|_{\square} \xrightarrow{\mathbb{P}} 0 \quad (40)$$

By Lemma A.9, it holds that

$$\delta_{\square}^{\text{bip}}(W_n, \mathcal{W}) = \delta_{\square}^{\text{bip}}(W_n^\pi, \mathcal{W}) \leq \|W_n^\pi - \mathcal{W}\|_{\square} \xrightarrow{\mathbb{P}} 0, \quad (41)$$

which completes the proof of Theorem 4.7.

B. Realisation of SNIP in the Factorised Saliency Model

In this appendix, we show that SNIP fits into the Factorised Saliency Model $S_{ij} \propto \phi_i \psi_j |\xi_{ij}|$ for a single hidden layer network.

B.1. Model and SNIP score at initialisation

Bounded input setting. Let the input be a single random sample

$$x = (x_1, \dots, x_d) \in \mathbb{R}^d, \quad x_i \stackrel{\text{i.i.d.}}{\sim} X, \quad |X| \leq B \text{ a.s.} \quad (42)$$

for some bound $B > 0$. Let $F_{|X|}$ denote the CDF of $|X|$ on $[0, B]$. We assume $F_{|X|}$ is continuous on $[0, B]$.

Initialisation. Let $(\theta_j^{(2)})_{j=1}^n$ and $(\theta_{ij}^{(1)})_{1 \leq i \leq d, 1 \leq j \leq n}$ be independent families with $\theta_j^{(2)} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$, $\theta_{ij}^{(1)} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$.

One hidden layer network. Consider the random network

$$f_n(x) = \frac{1}{\sqrt{n}} \sum_{j=1}^n \theta_j^{(2)} \sigma(h_j^{(1)}), \quad h_j^{(1)} := \frac{1}{\sqrt{d}} \sum_{i=1}^d \theta_{ij}^{(1)} x_i, \quad (43)$$

and squared loss on a fixed label $y \in \mathbb{R}$,

$$L = \frac{1}{2} (f_n(x) - y)^2, \quad \delta := f_n(x) - y. \quad (44)$$

SNIP pruning score. We introduce a binary mask $m_{ij}^{(1)} \in \{0, 1\}$ on first-layer weights via $\theta_{ij}^{(1)} \mapsto m_{ij}^{(1)} \theta_{ij}^{(1)}$ and evaluate SNIP at $m \equiv 1$. By the chain rule

$$\frac{\partial L}{\partial m_{ij}^{(1)}} = \delta \cdot \frac{\partial f_n(x)}{\partial m_{ij}^{(1)}} = \delta \cdot \frac{1}{\sqrt{n}} \theta_j^{(2)} \sigma'(h_j^{(1)}) \cdot \frac{1}{\sqrt{d}} \theta_{ij}^{(1)} x_i. \quad (45)$$

Hence, the SNIP pruning score magnitude factorises as

$$S_{ij}^{(n)} := \left| \frac{\partial L}{\partial m_{ij}^{(1)}} \right| = \frac{|\delta|}{\sqrt{nd}} \underbrace{|x_i|}_{\varphi_{n,i}} \underbrace{(|\theta_j^{(2)}| |\sigma'(h_j^{(1)})|)}_{\psi_{n,j}} \underbrace{|\theta_{ij}^{(1)}|}_{|\xi_{n,ij}|}. \quad (46)$$

Since top- ρ ranking is invariant to multiplying all scores by the same positive constant, we may equivalently work with the scale-free scores

$$S_{ij}^{e(n)} := \frac{\sqrt{nd}}{|\delta|} S_{ij}^{(n)} = \varphi_{n,i} \psi_{n,j} |\xi_{n,ij}| = |x_i| |\theta_j^{(2)}| |\sigma'(h_j^{(1)})| |\theta_{ij}^{(1)}|. \quad (47)$$

All masks below are defined using $S_{ij}^{e(n)}$.

B.2. Verification of Assumptions A1–A5 for the model

We verify the regularity assumptions in Section 4.2 for the SNIP factorisation $S_{ij}^{e(n)} = \varphi_{n,i} \psi_{n,j} |\xi_{n,ij}|$ with

$$\varphi_{n,i} = |x_i|, \quad \psi_{n,j} = |\theta_j^{(2)}| |\sigma'(h_j^{(1)})|, \quad \xi_{n,ij} = \theta_{ij}^{(1)}. \quad (48)$$

(A.4.1) Growth rate. Assumption A1 is a regime condition on (d, n) which we assume in advance.

(A.4.2) Deterministic input features via measure-preserving relabelling. Since $(|x_i|)_{i=1}^d$ are i.i.d. supported on $[0, B]$, we may relabel rows by sorting $\{|x_i|\}_{i=1}^d$ increasingly. Let $|x|_{(1)} \leq \dots \leq |x|_{(d)}$ be the order statistics and define $\varphi_{n,i} := |x|_{(i)}$. We define the deterministic limit profile $\varphi(u) := F_{|X|}^{-1}(u)$, $u \in [0, 1]$, which is bounded by B and continuous on $[0, 1]$ under continuity of $F_{|X|}$. Then standard uniform quantile consistency yields

$$\max_{1 \leq i \leq d} |\varphi_{n,i} - \varphi(i/d)| \xrightarrow{\mathbb{P}} 0 \quad (49)$$

This is exactly Assumption 4.2.

(A.4.3) Neuron features empirical CDF convergence Fix x and consider $(h_j^{(1)})_{j=1}^n$ with $h_j^{(1)} := \frac{1}{\sqrt{d}} \sum_{i=1}^d \theta_{ij}^{(1)} x_i$. Because the columns $(\theta_{ij}^{(1)})_{j=1}^n$ are i.i.d. and independent of x , we have:

- Conditional on x , the random variables $(h_j^{(1)})_{j=1}^n$ are i.i.d. across j as $d \rightarrow \infty$. Moreover, since $(\theta_{ij}^{(1)})_i$ are Gaussian,

$$h_j^{(1)} | x \sim \mathcal{N}(0, s_d^2(x)), \quad s_d^2(x) := \frac{\|x\|_2^2}{d}. \quad (50)$$

- $(\theta_j^{(2)})_{j=1}^n$ are i.i.d. and independent of $(h_j^{(1)})_{j=1}^n$. Therefore $\psi_{n,j} := |\theta_j^{(2)}| |\sigma'(h_j^{(1)})|$ are i.i.d. across j conditional on x .

To match Assumption 4.3 as stated (a non-random CDF F_ψ), we identify a deterministic limit law by showing $s_d^2(x) \rightarrow \nu^2$ with $\nu^2 := \mathbb{E}[X^2]$. Since $|X| \leq B$ a.s., Hoeffding's inequality implies for all $t > 0$,

$$\mathbb{P}\left(\left|\frac{1}{d} \sum_{i=1}^d x_i^2 - \nu^2\right| > t\right) \leq 2 \exp\left(-\frac{2dt^2}{B^4}\right), \quad (51)$$

hence $s_d^2(x) = \|x\|_2^2/d \rightarrow \nu^2$ in probability. Combining with Eqn. (50), the conditional law of $h_j^{(1)} | x$ converges to $\mathcal{N}(0, \nu^2)$, and thus $\psi_{n,j}$ converges in distribution to the deterministic limit

$$\psi \stackrel{d}{=} |\Theta^{(2)}| |\sigma'(\nu G)|, \quad \Theta^{(2)} \sim \mathcal{N}(0, 1), \quad G \sim \mathcal{N}(0, 1). \quad (52)$$

Let F_ψ denote the CDF of ψ on $[0, \infty)$ and quantile Q_ψ . Then the empirical CDF of $\{\psi_{n,j}\}_{j=1}^n$, $\hat{F}_{\psi,n}(t) := \frac{1}{n} \sum_{j=1}^n \mathbf{1}\{\psi_{n,j} \leq t\}$, satisfies $\hat{F}_{\psi,n}(t) \xrightarrow{\mathbb{P}} F_\psi(t)$, at continuity points t of F_ψ , Lebesgue-a.e. $v \in (0, 1)$.

(A.4.4) Edge noise and asymptotic decoupling. Let

$$\xi_{n,ij} := \theta_{ij}^{(1)} \sim \mathcal{N}(0, 1),$$

so that $|\xi|$ has a continuous distribution on $[0, \infty)$. The only minor subtlety is that

$$\psi_{n,j} = |\theta_j^{(2)}| |\sigma'(h_j^{(1)})|, \quad h_j^{(1)} = \frac{1}{\sqrt{d}} \sum_{r=1}^d \theta_{rj}^{(1)} x_r,$$

is not exactly independent of $\theta_{ij}^{(1)}$ at finite d , since $\theta_{ij}^{(1)}$ appears inside $h_j^{(1)}$. We remove this dependence by a leave-one-out replacement.

For each pair (i, j) , define

$$h_j^{(1),-i} := h_j^{(1)} - \frac{1}{\sqrt{d}} \theta_{ij}^{(1)} x_i = \frac{1}{\sqrt{d}} \sum_{r \neq i} \theta_{rj}^{(1)} x_r,$$

and

$$\psi_{n,j}^{-i} := |\theta_j^{(2)}| |\sigma'(h_j^{(1),-i})|.$$

Conditional on x , the variable $\psi_{n,j}^{-i}$ is independent of $\theta_{ij}^{(1)}$. Hence the proxy score

$$S_{ij}^{(n),-i} = |x_i| \psi_{n,j}^{-i} |\theta_{ij}^{(1)}|$$

has the desired factorised form with independent edge noise.

It remains to check that the leave-one-out replacement is asymptotically negligible. If σ' is bounded and Lipschitz, then

$$|\psi_{n,j} - \psi_{n,j}^{-i}| \leq |\theta_j^{(2)}| \text{Lip}(\sigma') \frac{|x_i| |\theta_{ij}^{(1)}|}{\sqrt{d}}.$$

Since the Gaussian weights have finite moments and $|x_i| \leq B$,

$$\mathbb{E} \left[\frac{1}{dn} \sum_{i=1}^d \sum_{j=1}^n |\psi_{n,j} - \psi_{n,j}^{-i}| \middle| x \right] \leq \frac{C}{\sqrt{d}}.$$

Therefore, by Markov's inequality,

$$\frac{1}{dn} \sum_{i=1}^d \sum_{j=1}^n |\psi_{n,j} - \psi_{n,j}^{-i}| \xrightarrow{\mathbb{P}} 0.$$

Thus the true SNIP scores are asymptotically equivalent to the leave-one-out scores in the sense needed for thresholded masks. Indeed, multiplying the above difference by the bounded row factor $|x_i|$ and the finite-moment edge factor $|\theta_{ij}^{(1)}|$ gives an average $o(1)$ perturbation of the score array. By the standard threshold-stability argument at continuity points of the limiting score distribution, the masks generated by $S_{ij}^{(n)}$ and $S_{ij}^{(n),-i}$ differ on a vanishing fraction of edges. Consequently, the graphon limit may be proved for the decoupled leave-one-out proxy and then transferred back to the original SNIP mask.

(A4.5) Threshold stability. The main theorem assumes the pruning thresholds τ_n converge to a deterministic limit $\tau \in (0, \infty)$. In practical PaI methods such as SNIP, sparsity is enforced by top- ρ selection, so the threshold is the empirical $(1 - \rho)$ -quantile $\hat{\tau}_n$ of the scores. Let \hat{F}_S denote the empirical CDF of the score array $\{S_{ij}\}_{i \leq d, j \leq n}$, and let F_S be the deterministic limiting CDF. As $(d, n) \rightarrow \infty$, \hat{F}_S concentrates uniformly around F_S , i.e., $\sup_{t \in \mathbb{R}} |\hat{F}_S(t) - F_S(t)| \xrightarrow{\mathbb{P}} 0$, which in our factorised setting follows from the conditional i.i.d. structure of $\{S_{ij}\}$ given the row and column features (A.4) together with a standard application of the Law of Large Number.

Define the theoretical threshold $\tau := F_S^{-1}(1 - \rho)$ and assume F_S is continuous at τ . Then $\hat{\tau}_n \rightarrow \tau$ in probability, and replacing $\hat{\tau}_n$ by the constant τ changes only a vanishing fraction of mask entries:

$$\frac{1}{dn} \sum_{i \leq d} \sum_{j \leq n} \left| \mathbf{1}\{S_{ij} > \hat{\tau}_n\} - \mathbf{1}\{S_{ij} > \tau\} \right| \xrightarrow{\mathbb{P}} 0. \quad (53)$$

Consequently, any graphon limit established for the deterministic-threshold mask $\mathbf{1}\{S_{ij} > \tau\}$ transfers directly to the practical top- ρ mask $\mathbf{1}\{S_{ij} > \hat{\tau}_n\}$.

B.3. Empirical Verification

We empirically illustrate Theorem 4.7 by comparing finite-width SNIP masks to the predicted limit graphon. For visualisation, we generate square masks $M_n \in \{0, 1\}^{n \times n}$ at two fixed densities $\rho = \{0.1, 0.2\}$ for $n \in \{200, 500, 1000, 2000, 4000\}$. Following Pham et al. (2025), we sort rows/columns by their latent factors (φ_i for rows and ψ_j for columns) and average the sorted masks over 100 random seeds to obtain an empirical edge-probability matrix.

Besides, we compute the limit theoretical graphon, $\mathcal{W}(u, v)$, as a continuous function on the unit square $[0, 1]^2$. To make the limit object computable, we consider i.i.d. standard normal inputs and weights initialisation. Since inputs $x \sim \mathcal{N}(0, 1)$, the magnitude $|x|$ follows a standard Half-Normal (HN) distribution. We define $\varphi(u) = F_{\text{HN}}^{-1}(u) = \sqrt{2} \operatorname{erf}^{-1}(u)$. The column factor is defined as $\psi = |\theta^{(2)}| |\sigma'(h)|$, where $\theta^{(2)} \sim \mathcal{N}(0, 1)$, and different activation functions σ (ReLU, Tanh, Sigmoid). We determine the quantile function $Q_\psi(v)$ for this product distribution via Monte Carlo simulation ($N = 10^6$ samples). The edge noise $|\xi_{ij}|$ is Half-Normal distributed with scale $\sigma_{\theta^{(1)}} = 1$ where $\theta^{(1)} \sim \mathcal{N}(0, 1)$. Given a global threshold τ (solved numerically to match density ρ), the theoretical edge probability is $\mathcal{W}_{(u,v)} = \mathbb{P}(\varphi(u)\psi(v)|\xi| > \tau) = 1 - \operatorname{erf}\left(\frac{\tau}{\varphi(u)Q_\psi(v)\sqrt{2}}\right)$.

Figure 3 and 4 shows that the averaged empirical masks become progressively smoother and rapidly approach the theoretical \mathcal{W} as n increases; by $n \geq 2000$ the empirical and theoretical patterns are visually indistinguishable. This supports the interpretation of PaI masks as finite-sample discretisations of a deterministic graphon limit.




Figure 3. **Visual Convergence to the Graphon Limit.** We compare empirical masks with different activation functions at increasing widths ($n = 200, 500, 1000, 2000, 4000$) against the analytically computed Theoretical Graphon. The density is fixed at $\rho = 0.1$.

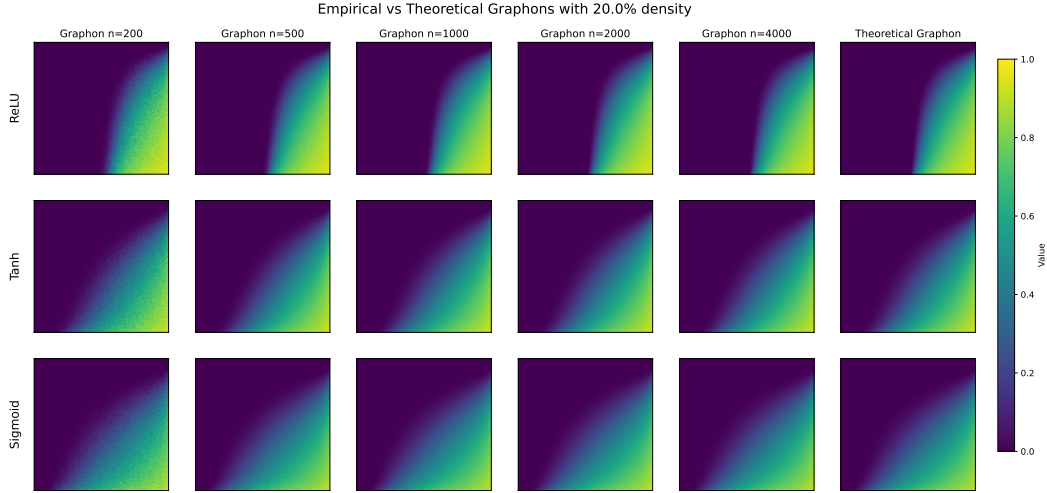


Figure 4. **Visual Convergence to the Graphon Limit.** We compare empirical masks with different activation functions at increasing widths ($n = 200, 500, 1000, 2000, 4000$) against the analytically computed Theoretical Graphon. The density is fixed at $\rho = 0.2$.

C. Realisation of GraSP in the Factorised Saliency Model

In this appendix, we show that GraSP can be realised into the Factorised Saliency Model by an explicit decomposition of the Hessian–gradient product in the same setting as in Appendix B. We treat both (i) the original signed GraSP pruning score and (ii) its magnitude variant, and show that, up to an (i, j) -independent scalar, they are asymptotically rank-equivalent to the corresponding signed or magnitude SNIP scores. The discrepancy is either identically zero for piecewise-linear activations, or vanishes uniformly as $n \rightarrow \infty$.

Recall that GraSP assigns each parameter a signed score based on a Hessian–gradient product. Let $\theta \in \mathbb{R}^{dn}$ be the vector of first-layer weights $\{\theta_{ij}^{(1)}\}_{1 \leq i \leq d, 1 \leq j \leq n}$, and define the gradient and Hessian of the loss w.r.t. θ as

$$s(\theta) := \nabla_{\theta} L(\theta) \in \mathbb{R}^{dn}, \quad H(\theta) := \nabla_{\theta}^2 L(\theta) \in \mathbb{R}^{dn \times dn}. \quad (54)$$

GraSP uses the elementwise product Hs to score coordinates. In matrix indexing, define the signed GraSP score as

$$Z_{ij}^{(n)} := -\theta_{ij}^{(1)} (Hs)_{ij}, \quad 1 \leq i \leq d, 1 \leq j \leq n. \quad (55)$$

Two pruning conventions. We consider:

1. **(Signed GraSP)**. Define the signed retention saliency

$$S_{ij}^{\text{sgn}} := -Z_{ij}^{(n)} = \theta_{ij}^{(1)}(Hs)_{ij}. \quad (56)$$

Then the signed GraSP mask is

$$M_{ij}^{\text{sgn}} := \mathbf{1}\{S_{ij}^{\text{sgn}} > \tau_n^{\text{sgn}}\}. \quad (57)$$

Note that τ_n^{sgn} may be negative. This variant retains weights based on the signed influence captured by Hs .

2. **(Magnitude GraSP)**. Define the magnitude saliency

$$S_{ij}^{\text{mag}} := |Z_{ij}^{(n)}|. \quad (58)$$

Then the magnitude GraSP mask is

$$M_{ij}^{\text{mag}} := \mathbf{1}\{S_{ij}^{\text{mag}} > \tau_n^{\text{mag}}\}, \quad \tau_n^{\text{mag}} \geq 0. \quad (59)$$

This variant retains weights with large influence regardless of sign and fits the nonnegative Factorised Saliency Model most directly.

Remark C.1. If masks are produced by top- ρ selection, then Eqn. (57) and Eqn. (59) correspond to taking τ^n to be the empirical $(1 - \rho)$ -quantile of $\{S_{ij}\}$. All results below apply to either form, via the same threshold-replacement argument as in Appendix B.4.

C.1. Hessian factorisation and the global scalar c_n

Let us write the Jacobian of the scalar output w.r.t. θ as $J := \nabla_{\theta} f_n(x) \in \mathbb{R}^{dn}$. By the chain rule, $g = \nabla_{\theta} L = \delta \nabla_{\theta} f_n = \delta J$. For squared loss, we have $H = \nabla_{\theta}^2 L = JJ^{\top} + \delta \nabla_{\theta}^2 f_n$. Therefore the Hessian–gradient product decomposes as

$$Hg = (JJ^{\top})(\delta J) + \delta(\nabla_{\theta}^2 f_n)(\delta J) = \underbrace{\|J\|_2^2}_{=:c_n} g + \underbrace{\delta^2(\nabla_{\theta}^2 f_n)J}_{=:R}. \quad (60)$$

The key point is that $c_n = \|J\|_2^2$ is a single scalar, which is independent of (i, j) .

Explicit formula for c_n . In our model,

$$J_{ij} = \frac{\partial f_n}{\partial \theta_{ij}^{(1)}} = \frac{1}{\sqrt{n}} \theta_j^{(2)} \sigma'(h_j^{(1)}) \cdot \frac{1}{\sqrt{d}} x_i, \quad (61)$$

so that

$$c_n = \sum_{i=1}^d \sum_{j=1}^n J_{ij}^2 = \left(\frac{\|x\|_2^2}{d}\right) \cdot \left(\frac{1}{n} \sum_{j=1}^n (\theta_j^{(2)})^2 \sigma'(h_j^{(1)})^2\right). \quad (62)$$

Explicit formula for the remainder R . A direct differentiation shows that the only nonzero second derivatives couple weights within the same neuron:

$$\frac{\partial^2 f_n}{\partial \theta_{ij}^{(1)} \partial \theta_{i'j'}^{(1)}} = \mathbf{1}\{j = j'\} \cdot \frac{1}{\sqrt{n}} \theta_j^{(2)} \sigma''(h_j^{(1)}) \cdot \frac{1}{d} x_i x_{i'}. \quad (63)$$

Using $J_{i'j} = \frac{1}{\sqrt{n}} \theta_j^{(2)} \sigma'(h_j^{(1)}) \frac{1}{\sqrt{d}} x_{i'}$ yields

$$R_{ij} = \delta^2 \sum_{i'=1}^d \frac{\partial^2 f_n}{\partial \theta_{ij}^{(1)} \partial \theta_{i'j}^{(1)}} J_{i'j} = \delta^2 \cdot \frac{(\theta_j^{(2)})^2}{n} \sigma''(h_j^{(1)}) \sigma'(h_j^{(1)}) \cdot \left(\frac{\|x\|_2^2}{d}\right) \cdot \frac{x_i}{\sqrt{d}}. \quad (64)$$

Combining Eqn. (55) and Eqn. (60) gives the entrywise score decomposition

$$Z_{ij}^{(n)} = -c_n \underbrace{(\theta_{ij}^{(1)} g_{ij})}_{\text{SNIP signed score}} - \underbrace{\theta_{ij}^{(1)} R_{ij}}_{=:E_{ij}}. \quad (65)$$

Recall from Appendix B that the SNIP score magnitude is $|\theta_{ij}^{(1)} g_{ij}| = \frac{|\delta|}{\sqrt{nd}} |x_i| |\theta_j^{(2)}| |\sigma'(h_j^{(1)})| |\theta_{ij}^{(1)}|$.

C.2. Lower-order remainder bounds

We consider two regimes.

Case 1: Piecewise-linear activations. If σ is piecewise linear (e.g. ReLU or leaky-ReLU), then $\sigma''(t) = 0$ for all $t \neq 0$. Since each $h_j^{(1)}$ has a continuous distribution at random initialization, $\mathbb{P}(h_j^{(1)} = 0) = 0$, hence $\sigma''(h_j^{(1)}) = 0$ for every j . By Eqn. (64),

$$R \equiv 0 \quad \text{a.s.} \quad \Rightarrow \quad Hg = c_n g. \quad (66)$$

Thus $Z_{ij}^{(n)} = -c_n(\theta_{ij}^{(1)} g_{ij})$ exactly, i.e. GraSP is an entrywise independent rescaling of SNIP.

Case 2: Smooth activations. Assume σ is twice differentiable with bounded second derivative $\|\sigma''\|_\infty := \sup_{t \in \mathbb{R}} |\sigma''(t)| < \infty$, and assume the nondegeneracy condition

$$\mu_1 := \mathbb{E}\left[(\Theta^{(2)})^2 \sigma'(\nu G)^2\right] > 0, \quad \Theta^{(2)}, G \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad \nu^2 := \mathbb{E}[X^2]. \quad (67)$$

Here ν^2 is the same deterministic variance limit used in the SNIP verification; see Appendix B.2.

Define $\bar{c}_n := \frac{1}{n} \sum_{j=1}^n (\theta_j^{(2)})^2 \sigma'(h_j^{(1)})^2$, so $c_n = (\|x\|_2^2/d) \bar{c}_n$ by Eqn. (62). Then $\bar{c}_n \rightarrow \mu_1$ in probability by the same argument used for the SNIP column features in Appendix B.2. In particular, for any $\eta \in (0, \mu_1)$,

$$\mathbb{P}(\bar{c}_n \geq \mu_1 - \eta) \rightarrow 1. \quad (68)$$

Next, we compare the remainder to the leading term. Using Eqn. (64) and $g_{ij} = \delta J_{ij} = \delta \cdot \frac{1}{\sqrt{n}} \theta_j^{(2)} \sigma'(h_j^{(1)}) \frac{1}{\sqrt{d}} x_i$, we obtain

$$\frac{|R_{ij}|}{|c_n g_{ij}|} = \frac{|\delta|}{\sqrt{n}} \cdot \frac{|\theta_j^{(2)}| |\sigma''(h_j^{(1)})|}{\bar{c}_n}. \quad (69)$$

Thus, on the event $\{\bar{c}_n \geq \mu_1/2\}$ and using $|\sigma''(h_j^{(1)})| \leq \|\sigma''\|_\infty$,

$$\max_{i \leq d, j \leq n} \frac{|R_{ij}|}{|c_n g_{ij}|} \leq \frac{2\|\sigma''\|_\infty}{\mu_1} \cdot \frac{|\delta|}{\sqrt{n}} \cdot \max_{j \leq n} |\theta_j^{(2)}|. \quad (70)$$

Finally, $\max_{j \leq n} |\theta_j^{(2)}| = O(\sqrt{\log n})$ by Gaussian maxima, and $\delta = f_n(x) - y = O(1)$ at initialization under the $1/\sqrt{n}$ output scaling and bounded-input regime, as $f_n(x)$ is sub-Gaussian conditioned on x . Therefore the right-hand side of Eqn. (70) is $O(\sqrt{\log n/n}) \rightarrow 0$. Combining with Eqn. (68), we conclude the uniform vanishing:

$$\max_{i \leq d, j \leq n} \frac{|R_{ij}|}{|c_n g_{ij}|} \xrightarrow{\mathbb{P}} 0 \quad \text{as } n \rightarrow \infty. \quad (71)$$

Consequence: Asymptotic rank-equivalence to SNIP. By Eqn. (65) and Eqn. (71),

$$Z_{ij}^{(n)} = -c_n(\theta_{ij}^{(1)} g_{ij}) \cdot (1 + o(1)) \quad \text{uniformly over the entries.}$$

Therefore the signed score $S_{ij}^{\text{sgn}} = -Z_{ij}^{(n)}$ is uniformly rank-equivalent to the signed SNIP score $\theta_{ij}^{(1)} s_{ij}$, while the magnitude saliency $S_{ij}^{\text{mag}} = |Z_{ij}^{(n)}|$ is uniformly rank-equivalent to the magnitude SNIP score $|\theta_{ij}^{(1)} s_{ij}|$.

C.3. Realising GraSP in the factorised saliency model and assumption verification

The leading order GraSP score depends on (i, j) only through the factor $\theta_{ij}^{(1)} g_{ij}$ up to the global scalar c_n . In particular, the leading term factorises exactly as in the SNIP appendix:

$$\theta_{ij}^{(1)} g_{ij} = \frac{\delta}{\sqrt{nd}} x_i (\theta_j^{(2)} \sigma'(h_j^{(1)})) \theta_{ij}^{(1)}, \quad (72)$$

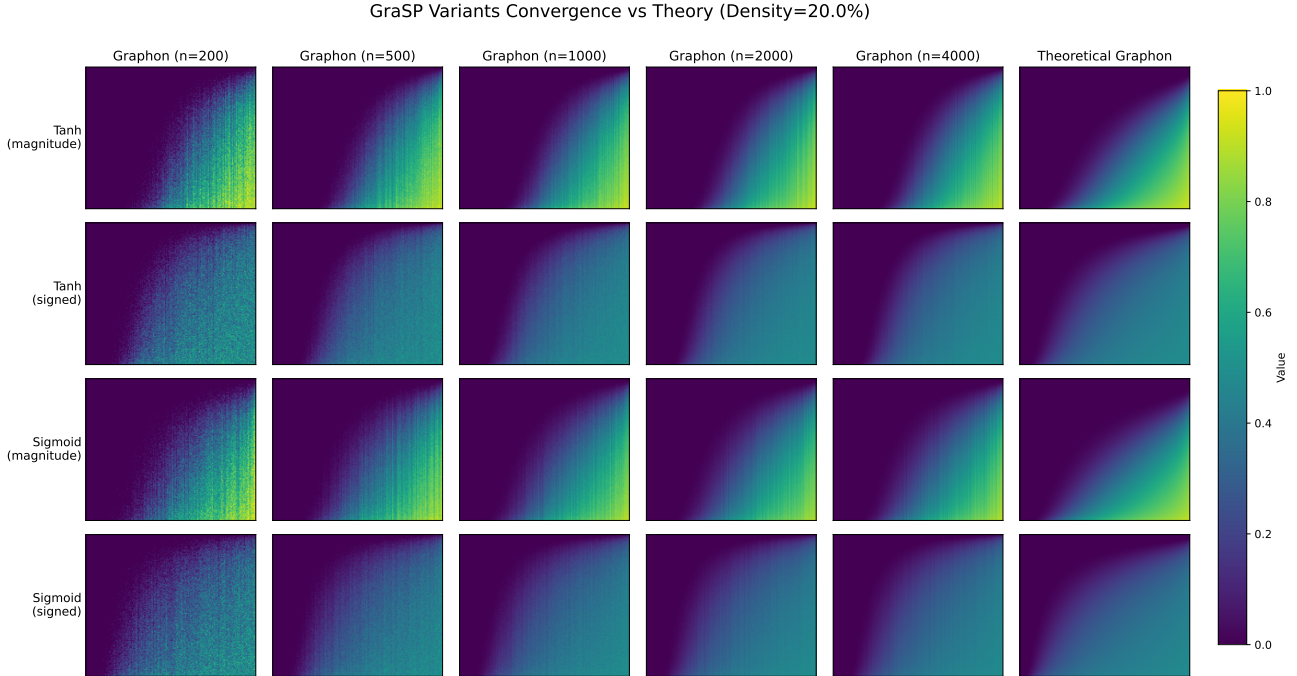


Figure 5. **Graphon Convergence of GraSP Variants.** Empirical pruning masks for GraSP at increasing widths ($n \in \{200, \dots, 4000\}$) with density $\rho = 20\%$. We compare **Magnitude GraSP** (rows 1, 3) and **Signed GraSP** (rows 2, 4) for Tanh and Sigmoid activations.

Thus, the signed GraSP fits a *signed* factorised model

$$S_{n,ij}^{\text{sgn}} = \varphi_{n,i} \psi_{n,j} \xi_{n,ij}, \quad M_{n,ij}^{\text{sgn}} = \mathbf{1}\{S_{n,ij}^{\text{sgn}} > \tau_n^{\text{sgn}}\}, \quad (73)$$

with $\varphi_{n,i} = x_i$, $\psi_{n,j} = \theta_j^{(2)} \sigma'(h_j^{(1)})$, $\xi_{n,ij} = \theta_{ij}^{(1)}$.

Meanwhile, with magnitude GraSP, we take absolute values yields

$$|\theta_{ij}^{(1)} s_{ij}| \propto |x_i| (|\theta_j^{(2)}| |\sigma'(h_j^{(1)})|) |\theta_{ij}^{(1)}|, \quad (74)$$

which is exactly of the non-negative factorised saliency model form:

$$S_{n,ij}^{\text{mag}} = \varphi_{n,i} \psi_{n,j} |\xi_{n,ij}|, \quad M_{n,ij}^{\text{mag}} = \mathbf{1}\{S_{n,ij}^{\text{mag}} > \tau_n^{\text{mag}}\}, \quad \tau_n^{\text{mag}} \geq 0. \quad (75)$$

Therefore, the verification of Assumptions 4.1–4.5 for GraSP follows the same template as SNIP (Appendix B.2): row-profile quantile consistency for $\varphi_{n,i}$, empirical CDF convergence for $\psi_{n,j}$, asymptotic decoupling of $\xi_{n,ij}$ from $\psi_{n,j}$ under bounded inputs, and threshold stability/threshold replacement for top- ρ pruning (Appendix B.2–B.4). The only additional ingredient for GraSP is the remainder control Eqn. (71), ensuring that the Hessian correction does not alter the limiting ranking and mask.

C.4. Empirical verification and theoretical graphon details

We empirically validate the convergence of GraSP to its theoretical limit using the setup described in Appendix B. We vary network width $n \in \{200, \dots, 4000\}$ and fix density $\rho = 0.2$. We exclude ReLU activations here, as the second derivative vanishes ($\sigma'' \equiv 0$) almost everywhere, rendering the Hessian term zero in standard implementations; instead, we focus on Tanh and Sigmoid to verify the non-trivial asymptotic equivalence in the smooth regime. Figure 5 presents the results. To visualise the latent structure, rows and columns are sorted by the theoretical latent factors $|\varphi(u)|$ and $\psi(v)$ respectively.

Convergence to the limit. For both activations, the empirical masks smooth out as n increases, becoming visually indistinguishable from the theoretical prediction at $n = 4000$. This confirms that despite the complex Hessian interactions, the GraSP mask converges to a deterministic graphon.

Theoretical Graphon Details: While our derivation in Appendix C.2 proves that Signed and Magnitude GraSP each is rank-equivalent to its corresponding SNIP variant; the two induce different link functions, Figure 5 reveals a distinct visual difference: the *Signed* variant (rows 2 and 4) exhibits a much wider transition boundary than the *Magnitude* variant. This difference is explained by the specific link functions induced by their respective selection rules. Let $\varphi(u)\psi(v)$ be the latent signal strength and $\xi \sim \mathcal{N}(0, 1)$ be the edge noise.

- **Magnitude GraSP.** This variant retains edges where $|\varphi(u)\psi(v) \cdot \xi| > \tau^{\text{mag}}$. The limiting edge probability is:

$$W^{\text{mag}}(u, v) = \mathbb{P} \left(|\xi| > \frac{\tau^{\text{mag}}}{\varphi(u)\psi(v)} \right) = 1 - \text{erf} \left(\frac{\tau^{\text{mag}}}{\varphi(u)\psi(v)\sqrt{2}} \right). \quad (76)$$

The error function decays exponentially fast. Furthermore, to target a sparsity of $\rho = 0.2$, the threshold τ^{mag} must be relatively large (far into the tails), resulting in a fast transition from $W \approx 1$ to $W \approx 0$ as $\varphi(u)\psi(v)$ decreases.

- **Signed GraSP.** This variant retains edges where $\varphi(u)\psi(v) \cdot \xi > \tau^{\text{sgn}}$. The limiting edge probability is:

$$W^{\text{sgn}}(u, v) = \mathbb{P} \left(\xi > \frac{\tau^{\text{sgn}}}{\varphi(u)\psi(v)} \right) = 1 - \Phi \left(\frac{\tau^{\text{sgn}}}{\varphi(u)\psi(v)} \right), \quad (77)$$

where Φ is the Standard Normal CDF. To achieve the *same* density (i.e., $\rho = 0.2$) using only the positive tail, the signed threshold τ^{sgn} must be significantly lower (closer to the mean) than τ^{mag} . Operating closer to the probability mass center means that small changes in the signal $\varphi(u)\psi(v)$ result in more gradual changes in edge probability.

The rightmost column of Figure 5 plots these exact functions. The agreement between the $n = 4000$ empirical masks and their respective theoretical counterparts confirms that the wider spread in Signed GraSP is not an artifact of slow convergence, but the mathematically correct limiting behavior of a one-sided selection rule.

D. Extension of Graphon Convergence Theorem to Deep Neural Network

Consider an L -hidden layer network at initialisation with widths $n_0 = d, n_1, \dots, n_L$ and forward recursion

$$h_j^{(\ell)} := \frac{1}{\sqrt{n_{\ell-1}}} \sum_{i=1}^{n_{\ell-1}} \theta_{ij}^{(\ell)} z_i^{(\ell-1)}, \quad z_j^{(\ell)} := \sigma(h_j^{(\ell)}), \quad \ell = 1, \dots, L, \quad (78)$$

where $\{\theta_{ij}^{(\ell)}\}$ are i.i.d. across entries and independent across layers, and σ is an activation function.

In the deep factorised saliency model at layer ℓ ,

$$S_{n,ij}^{(\ell)} = \varphi_{n,i}^{(\ell)} \psi_{n,j}^{(\ell)} |\xi_{n,ij}^{(\ell)}|, \quad M_{n,ij}^{(\ell)} = \mathbf{1}\{S_{n,ij}^{(\ell)} > \tau_n^{(\ell)}\}, \quad i \in [n_{\ell-1}], j \in [n_\ell], \quad (79)$$

Now, for SNIP saliency score, we show that for a hidden layer ℓ , Assumptions A4.1 - A4.5 are satisfied asymptotically.

Deterministic input feature limit. Fix a hidden layer $\ell \geq 2$. For the layer- ℓ mask, the input feature is the magnitude of the previous-layer activation,

$$\varphi_{n,i}^{(\ell)} = |z_i^{(\ell-1)}|, \quad i \in [n_{\ell-1}]. \quad (80)$$

We verify that the empirical distribution of $\{\varphi_{n,i}^{(\ell)}\}_{i=1}^{n_{\ell-1}}$ has a deterministic limit. Conditional on $z^{(\ell-2)}$, the pre-activations in layer $\ell - 1$ satisfy

$$h_i^{(\ell-1)} = \frac{1}{\sqrt{n_{\ell-2}}} \sum_{r=1}^{n_{\ell-2}} \theta_{ri}^{(\ell-1)} z_r^{(\ell-2)}, \quad i \in [n_{\ell-1}]. \quad (81)$$

Since the columns $\{\theta_{:,i}^{(\ell-1)}\}_{i=1}^{n_{\ell-1}}$ are i.i.d., the random variables $\{h_i^{(\ell-1)}\}_{i=1}^{n_{\ell-1}}$ are i.i.d. conditional on $z^{(\ell-2)}$. In the Gaussian initialisation case,

$$h_i^{(\ell-1)} \mid z^{(\ell-2)} \sim \mathcal{N}\left(0, q_n^{(\ell-2)}\right), \quad q_n^{(\ell-2)} := \frac{1}{n_{\ell-2}} \sum_{r=1}^{n_{\ell-2}} (z_r^{(\ell-2)})^2. \quad (82)$$

Under the standard infinite-width moment assumptions, the empirical second moment satisfies

$$q_n^{(\ell-2)} \xrightarrow{\mathbb{P}} q^{(\ell-2)}, \quad (83)$$

where $q^{(\ell-2)}$ is deterministic. Hence, the conditional law of $h_i^{(\ell-1)}$ converges to the deterministic Gaussian law $\mathcal{N}(0, q^{(\ell-2)})$.

Therefore the row feature $\varphi_{n,i}^{(\ell)} = |z_i^{(\ell-1)}|$ has the deterministic limiting law $\varphi^{(\ell)} = |\sigma(h^{(\ell-1)})|$. Let $F_{\varphi^{(\ell)}}$ denote the CDF of $\varphi^{(\ell)}$. The empirical CDF

$$\widehat{F}_{\varphi^{(\ell)}, n}(t) := \frac{1}{n_{\ell-1}} \sum_{i=1}^{n_{\ell-1}} \mathbf{1}\{\varphi_{n,i}^{(\ell)} \leq t\} \quad (84)$$

then satisfies

$$\sup_{t \in \mathbb{R}} \left| \widehat{F}_{\varphi^{(\ell)}, n}(t) - F_{\varphi^{(\ell)}}(t) \right| \xrightarrow{\mathbb{P}} 0. \quad (85)$$

Consequently, after sorting the row features increasingly, the empirical quantiles converge to the deterministic quantile curve

$$Q_{\varphi^{(\ell)}}(u) := \inf\{s \geq 0 : F_{\varphi^{(\ell)}}(s) \geq u\}. \quad (86)$$

In particular, at continuity points of $Q_{\varphi^{(\ell)}}$,

$$\varphi_{n, \lceil n_{\ell-1} u \rceil}^{(\ell)} \xrightarrow{\mathbb{P}} Q_{\varphi^{(\ell)}}(u), \quad u \in (0, 1). \quad (87)$$

This verifies the row-feature analogue of Assumption 4.2 for layer ℓ .

Neuron-feature empirical CDF convergence. For hidden layers, the neuron feature should be defined through the normalised backpropagated sensitivity rather than through the raw outgoing weight norm. We have $\psi_j^{(\ell)} := |\sigma'(h_j^{(\ell)}) b_j^{(\ell)}|$, where $b_j^{(L)} = \theta_j^{(L+1)}$ for the last hidden layer, and recursively, $b_j^{(\ell)} = \frac{1}{\sqrt{n_\ell}} \sum_{k=1}^{n_{\ell+1}} \theta_{jk}^{(\ell+1)} \sigma'(h_k^{(\ell+1)}) b_k^{(\ell+1)}$.

Under the standard infinite-width moment assumptions, the forward variables satisfy $h_j^{(\ell)} \sim \mathcal{N}(0, q_{\ell-1})$ with deterministic variance recursion $q_\ell = \mathbb{E}[\sigma(h^{(\ell)})^2]$. The backward factors also have a deterministic limiting law. Starting from $b_j^{(L)} \sim \mathcal{N}(0, 1)$, define recursively

$$b_j^{(\ell)} \sim \mathcal{N}(0, s_{\ell+1}), \quad s_{\ell+1} = \mathbb{E}\left[(\sigma'(h^{(\ell+1)}) b^{(\ell+1)})^2\right],$$

with $b^{(\ell)}$ independent of $h^{(\ell)}$ in the infinite-width limit. Hence, $\psi^{(\ell)}$ has a deterministic limiting distribution. Let $F_{\psi^{(\ell)}}$ be its CDF. Then the empirical CDF $\widehat{F}_{\psi^{(\ell)}, n}(t) = \frac{1}{n_\ell} \sum_{j=1}^{n_\ell} \mathbf{1}\{\psi_{n,j}^{(\ell)} \leq t\}$ satisfies $\sup_{t \in \mathbb{R}} \left| \widehat{F}_{\psi^{(\ell)}, n}(t) - F_{\psi^{(\ell)}}(t) \right| \xrightarrow{\mathbb{P}} 0$.

This verifies the layerwise analogue of Assumption 4.3.

Edge noise asymptotic decoupling. The subtle point is that the neuron feature

$$\psi_j^{(\ell)} = |\sigma'(h_j^{(\ell)}) b_j^{(\ell)}| \tag{88}$$

is not exactly independent of the edge noise $\theta_{ij}^{(\ell)}$. This dependence enters directly through $h_j^{(\ell)}$, and indirectly through the downstream factor $b_j^{(\ell)}$, since perturbing $h_j^{(\ell)}$ changes later-layer activations and hence the backpropagated sensitivities.

For each edge (i, j) , define the leave-one-out pre-activation

$$h_j^{(\ell), -i} = h_j^{(\ell)} - \frac{1}{\sqrt{n_{\ell-1}}} \theta_{ij}^{(\ell)} z_i^{(\ell-1)}. \tag{89}$$

Let $b_j^{(\ell), -i}$ be the corresponding downstream factor computed in the leave-one-out network, and define

$$\psi_j^{(\ell), -i} = |\sigma'(h_j^{(\ell), -i}) b_j^{(\ell), -i}|. \tag{90}$$

Conditional on all weights and activations except $\theta_{ij}^{(\ell)}$, the leave-one-out feature $\psi_j^{(\ell), -i}$ is independent of $\theta_{ij}^{(\ell)}$.

Moreover, by Lipschitzness of σ' ,

$$|\psi_j^{(\ell)} - \psi_j^{(\ell), -i}| \leq \frac{L_{\sigma'}}{\sqrt{n_{\ell-1}}} |b_j^{(\ell)}| |\theta_{ij}^{(\ell)}| |z_i^{(\ell-1)}| + \|\sigma'\|_\infty |b_j^{(\ell)} - b_j^{(\ell), -i}|. \tag{91}$$

The first term is $O(n_{\ell-1}^{-1/2})$ after averaging over (i, j) , using the uniform moment bounds on $b_j^{(\ell)}$, $\theta_{ij}^{(\ell)}$, and $z_i^{(\ell-1)}$. The second term is controlled by the standard forward-backward stability recursion: the perturbation created by removing one edge has normalised size $O((n_{\ell-1} n_\ell)^{-1/2})$ at layer ℓ , remains of this order as it propagates forward through the fixed-depth network, and therefore induces an averaged backward-factor perturbation

$$\frac{1}{n_{\ell-1} n_\ell} \sum_{i=1}^{n_{\ell-1}} \sum_{j=1}^{n_\ell} |b_j^{(\ell)} - b_j^{(\ell), -i}| = O(n_{\ell-1}^{-1/2}). \tag{92}$$

Consequently,

$$\frac{1}{n_{\ell-1} n_\ell} \sum_{i=1}^{n_{\ell-1}} \sum_{j=1}^{n_\ell} |\psi_j^{(\ell)} - \psi_j^{(\ell), -i}| \xrightarrow{\mathbb{P}} 0. \tag{93}$$

Thus the original layer- ℓ scores and their leave-one-out decoupled versions differ by a vanishing average perturbation. At continuity points of the limiting score distribution, the corresponding thresholded masks differ on a vanishing fraction of entries, so their step-kernels differ by $o(1)$ in L^1 , and hence also in cut norm.

Threshold stability. For top- ρ pruning, let $\tau_n^{(\ell)}$ be the empirical $(1 - \rho)$ -quantile of the layer- ℓ scores. The layerwise empirical score distribution converges to a deterministic limit distribution. Assuming this limit distribution is continuous at its $(1 - \rho)$ -quantile, we obtain $\tau_n^{(\ell)} \xrightarrow{\mathbb{P}} \tau^{(\ell)}$. Replacing $\tau_n^{(\ell)}$ by $\tau^{(\ell)}$ changes only a vanishing fraction of mask entries.

Corollary D.1 (Layerwise graphon convergence in deep neural networks). *Fix a depth $L \in \mathbb{N}$. For each layer $\ell = 1, \dots, L$, let $M_n^{(\ell)} \in \{0, 1\}^{n_{\ell-1} \times n_{\ell}}$ be the PAI mask at layer ℓ generated by a factorised saliency model*

$$S_{n,ij}^{(\ell)} = \varphi_{n,i}^{(\ell)} \psi_{n,j}^{(\ell)} |\xi_{n,ij}^{(\ell)}|, \quad M_{n,ij}^{(\ell)} = \mathbf{1}\{S_{n,ij}^{(\ell)} > \tau_n^{(\ell)}\}. \quad (94)$$

Let $W_n^{(\ell)} := W_{M_n^{(\ell)}}$ be the associated bipartite step-kernel. Assume that for each fixed ℓ , the layerwise analogues of Assumptions A4.1 - A4.5 hold, with deterministic limiting graphon $\mathcal{W}^{(\ell)} \in L^\infty([0, 1]^2)$. Then, as $\min_{\ell} n_{\ell} \rightarrow \infty$,

$$\max_{\ell \in [L]} \delta_{\square}^{\text{bip}}(W_n^{(\ell)}, \mathcal{W}^{(\ell)}) \xrightarrow{\mathbb{P}} 0. \quad (95)$$

Proof sketch. Fix $\ell \in [L]$. By the layerwise Assumptions A4.1 - A4.5, Theorem 4.7 applies to the bipartite mask matrix $M_n^{(\ell)}$ and yields $\delta_{\square}^{\text{bip}}(W_n^{(\ell)}, \mathcal{W}^{(\ell)}) \rightarrow 0$ in probability. Moreover, in the deep setting the *input-feature* sequence at layer ℓ can be taken as $\varphi_{n,i}^{(\ell)} = z_i^{(\ell-1)}$, and as shown above the post-activations $\{z_i^{(\ell-1)}\}_{i=1}^{n_{\ell-1}}$ are (asymptotically) i.i.d. with a deterministic limit law. Hence, $\{\varphi_{n,i}^{(\ell)}\}$ satisfies the required input-feature assumption. Finally, since L is fixed, a union bound gives for any $\varepsilon > 0$,

$$\mathbb{P}\left(\max_{\ell \leq L} \delta_{\square}^{\text{bip}}(W_n^{(\ell)}, \mathcal{W}^{(\ell)}) > \varepsilon\right) \leq \sum_{\ell=1}^L \mathbb{P}\left(\delta_{\square}^{\text{bip}}(W_n^{(\ell)}, \mathcal{W}^{(\ell)}) > \varepsilon\right) \rightarrow 0, \quad (96)$$

□

D.1. Empirical Verification for 2-Hidden Layer Neural Networks

We extend our empirical analysis beyond the single-hidden-layer setting to verify the Graphon Limit Hypothesis in deep architectures. We examine the convergence behavior of a 2-hidden-layer fully connected network ($d \rightarrow n \rightarrow n \rightarrow 1$) under varying sparsity levels $\rho \in \{0.1, 0.2\}$ and activation functions $\sigma \in \{\text{ReLU}, \text{Tanh}, \text{Sigmoid}\}$.

Experimental Setup. Inputs are drawn from a standard normal distribution $x \sim \mathcal{N}(0, I_d)$. We generate SNIP masks for networks of increasing width $n \in \{200, 500, 1000, 2000, 4000\}$. To visualise the underlying graphon structure, the discrete adjacency matrices are sorted according to their latent factors:

- Layer 1 ($\mathcal{W}^{(1)}$): Rows are sorted by the input magnitude $|x_i|$.
- Layer 2 ($\mathcal{W}^{(2)}$): Rows are sorted by the magnitude of the incoming activations $|\sigma(h_j^{(1)})|$, where $h^{(1)}$ are the pre-activations of the first layer.

The results are averaged over 20 random seeds to smooth out finite-sample noise.

Theoretical Limit Construction. While the limit for layer 1 follows the Half-Normal profile derived in Section 4, the limit for layer 2 is derived recursively. As $n \rightarrow \infty$, the pre-activations $h^{(1)}$ converge to a Gaussian process with distribution $\mathcal{N}(0, 1)$ (up to variance scaling). Consequently, the theoretical row factors for layer 2 are modelled by the quantile function of $|\sigma(Z)|$ where $Z \sim \mathcal{N}(0, 1)$.

Results and Discussion. Figures 6 - 11 illustrate the convergence of empirical masks to the theoretical limits. Across all activations and densities, the empirical masks converge in the cut-metric topology to the predicted deterministic graphons. At $n = 4000$, the empirical masks are visually indistinguishable from the theoretical limit.

The choice of activation function significantly affects the topology of the graphon. For ReLU activations, since weights are initialised follow Normal distribution, the pre-activation h is symmetric around 0, leading to roughly half of the columns in layer 1 graphon being pruned. Then, in the second layer, both the inputs and the gradients are subject to the ReLU cutoff. This explains why it looks like three-quarters of the graphon are pruned in layer 2. In contrast, for activations like Tanh and Sigmoid, the high-density region in layer 2 appears more diffuse.

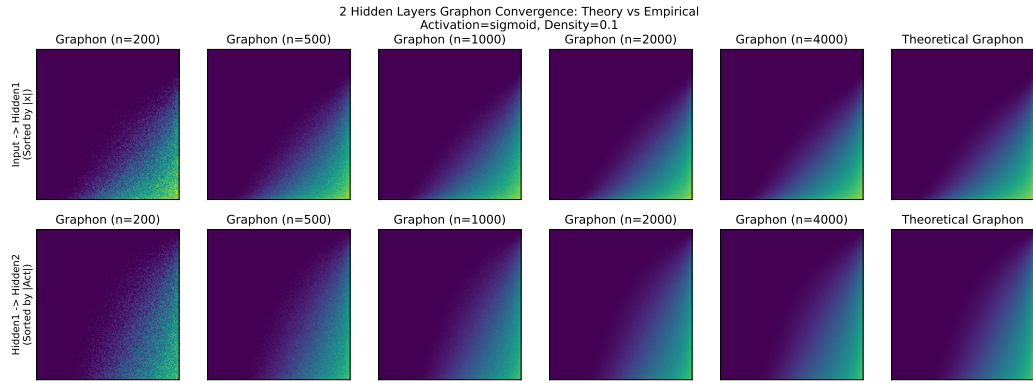


Figure 6. Graphon convergence in 2-hidden layer networks with sigmoid activation function on two densities $\rho = 0.1$.

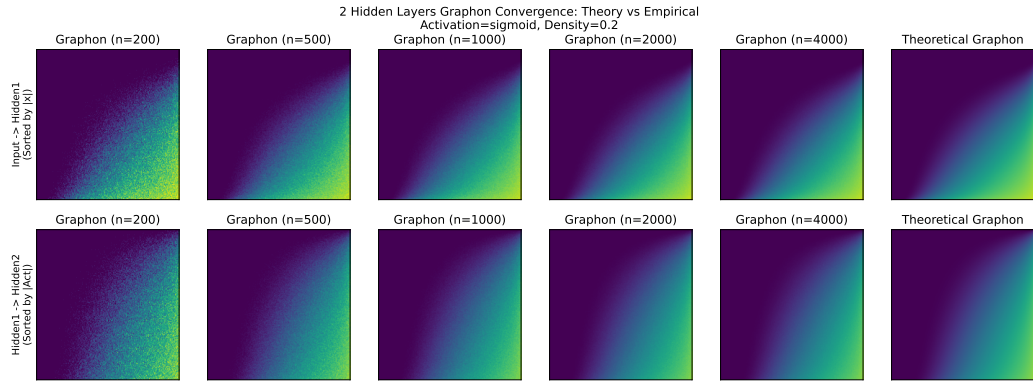


Figure 7. Graphon convergence in 2-hidden layer networks with sigmoid activation function on two densities $\rho = 0.2$.

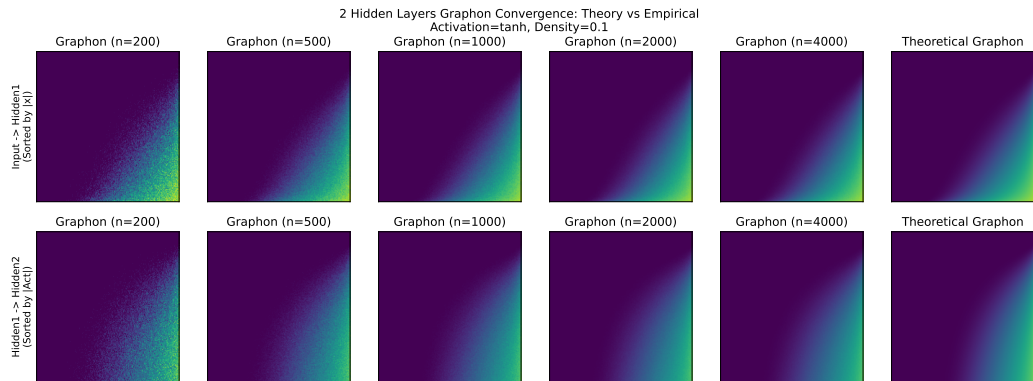


Figure 8. Graphon convergence in 2-hidden layer networks with tanh activation function on two densities $\rho = 0.1$.

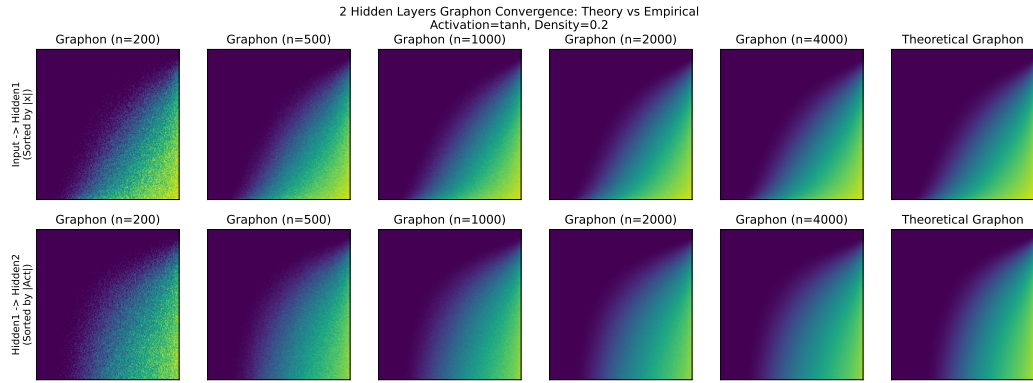


Figure 9. Graphon convergence in 2-hidden layer networks with tanh activation function on two densities $\rho = 0.2$.

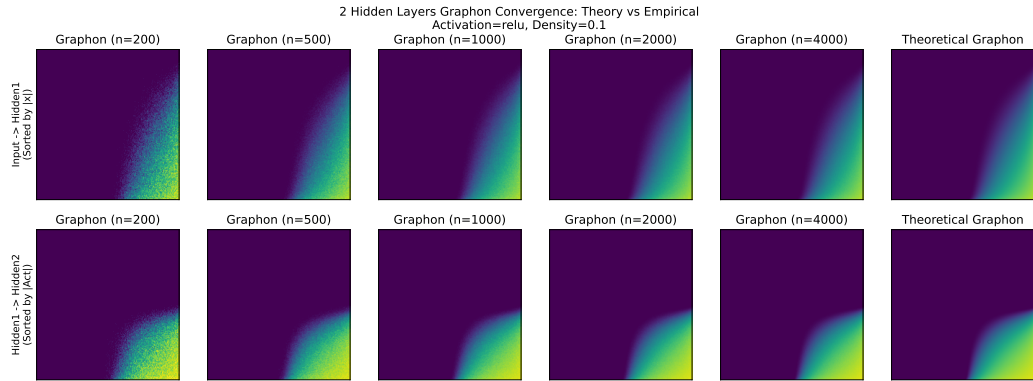


Figure 10. Graphon convergence in 2-hidden layer networks with ReLU activation function on two densities $\rho = 0.1$.

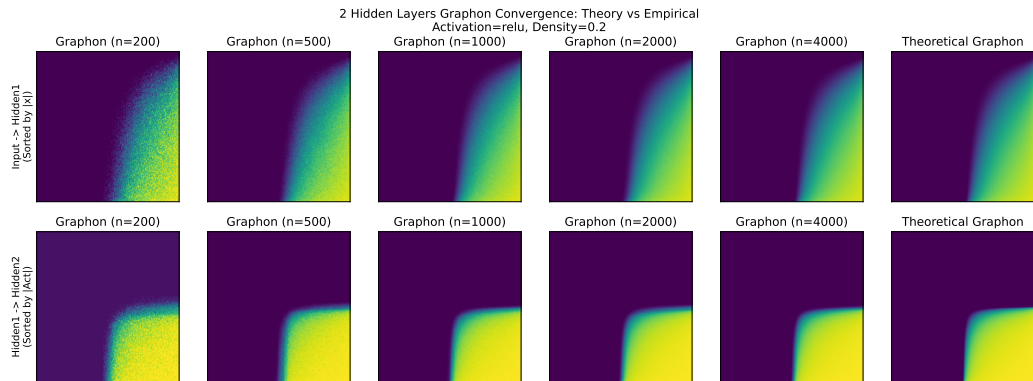


Figure 11. Graphon convergence in 2-hidden layer networks with ReLU activation function on two densities $\rho = 0.2$.

E. Proofs of Results in Section 5

In this appendix, we prove Lemma 5.6 and Theorem 5.7. Throughout, the mask is defined by $M_{ij} = \mathbf{1}\{\varphi_i \psi_j |\xi_{ij}| > \tau\}$, and conditioned on (φ, ψ) the noises $(\xi_{ij})_{i \leq d, j \leq n}$ are independent with common law P_ξ . Recall the survival link $g(\tau, z) := \mathbb{P}(z|\xi| > \tau)$ so that

$$\mathbb{P}(M_{ij} = 1 \mid \varphi, \psi) = g(\tau, \varphi_i \psi_j). \quad (97)$$

We work in the sorted coordinates $u_i = i/d$ and $v_j = j/n$ where $\max_{i \leq d} |\varphi_i - \varphi(u_i)| \rightarrow 0$ and $\max_{j \leq n} |\psi_j - Q_\psi(v_j)| \rightarrow 0$, and the limiting edge-probability kernel is

$$\mathcal{W}(u, v) = g(\tau, \varphi(u)Q_\psi(v)). \quad (98)$$

E.1. A counting argument for the active column set

Let $V_\star \subset (0, 1)$ be as in Assumption 5.2 with $\lambda(V_\star) = \alpha > 0$. Define the set of *active columns*

$$J_n := \{j \in [n] : v_j = j/n \in V_\star\}. \quad (99)$$

Since $v_j = j/n$ is deterministic, we have $|J_n|/n \rightarrow \alpha$ provided V_\star is Riemann measurable. that is ∂V_\star has Lebesgue measure 0, or V_\star is a finite union of intervals. In particular, for all sufficiently large n , it holds that

$$|J_n| \geq \frac{\alpha}{2} n. \quad (100)$$

E.2. Proof of Lemma 5.6

Proof of Lemma 5.6. Fix $\tilde{n} \in \mathbb{N}$ and let $k := |I|$. Let J_n be the deterministic active column set defined in Eqn. (99). First, we show a uniform edge-probability lower bound on the active rectangle. By Assumption 5.2, there exists $p_\star > 0$ such that $\mathcal{W}(u, v) \geq p_\star$ for a.e. $(u, v) \in U_\star \times V_\star$. Since $g(\tau, \cdot)$ is nondecreasing and continuous on bounded intervals, and since $\varphi_i \rightarrow \varphi(u_i)$ and $\psi_j \rightarrow Q_\psi(v_j)$ uniformly, we can choose n large enough so that on an event E_n with $\mathbb{P}(E_n) \rightarrow 1$,

$$\min_{i \in I} \min_{j \in J_n} g(\tau, \varphi_i \psi_j) \geq \frac{p_\star}{2}. \quad (101)$$

Secondly, we define good columns and compute their success probabilities. For each $j \in J_n$, define the indicator that column j is fully connected to the k active inputs:

$$G_j := \mathbf{1}\{M_{ij} = 1 \text{ for all } i \in I\}. \quad (102)$$

Let $\mathcal{G} := \sigma(\varphi, \psi)$ be the sigma-field generated by the row or column features. Conditional on \mathcal{G} , the edge noises are independent across (i, j) , hence the events $\{M_{ij} = 1\}_{i \in I}$ are independent for each fixed j . Therefore,

$$\mathbb{P}(G_j = 1 \mid \mathcal{G}) = \prod_{i \in I} \mathbb{P}(M_{ij} = 1 \mid \mathcal{G}) = \prod_{i \in I} g(\tau, \varphi_i \psi_j). \quad (103)$$

On the event E_n from Eqn. (101), we obtain the lower bound

$$\mathbb{P}(G_j = 1 \mid \mathcal{G}) \geq \left(\frac{p_\star}{2}\right)^k \quad \text{for all } j \in J_n. \quad (104)$$

Moreover, conditional on \mathcal{G} , the random variables $(G_j)_{j \in J_n}$ are independent: indeed, G_j depends only on the noises $\{\xi_{ij} : i \in I\}$, which are disjoint across distinct columns j .

Next, we use Chernoff bound for the number of good columns. Particularly, we define the number of good columns

$$N_{\text{good}} := \sum_{j \in J_n} G_j. \quad (105)$$

Conditionws on \mathcal{G} and on E_n , the $(G_j)_{j \in J_n}$ are independent Bernoulli variables with success probabilities bounded below as in Eqn. (104). Hence the conditional mean satisfies

$$\mu_n := \mathbb{E}[N_{\text{good}} \mid \mathcal{G}] = \sum_{j \in J_n} \mathbb{P}(G_j = 1 \mid \mathcal{G}) \geq |J_n| \left(\frac{p_\star}{2}\right)^k. \quad (106)$$

Using the deterministic bound Eqn. (100), we obtain on E_n :

$$\mu_n \geq \frac{\alpha}{2} n \left(\frac{p_\star}{2} \right)^k. \quad (107)$$

We now apply a standard multiplicative Chernoff bound for sums of independent Bernoulli variables: for any $\delta \in (0, 1)$,

$$\mathbb{P}\left(N_{\text{good}} \leq (1 - \delta)\mu_n \mid \mathcal{G}\right) \leq \exp\left(-\frac{\delta^2}{2}\mu_n\right). \quad (108)$$

Take $\delta = 1/2$. Then on E_n ,

$$\mathbb{P}\left(N_{\text{good}} \leq \frac{1}{2}\mu_n \mid \mathcal{G}\right) \leq \exp(-\mu_n/8). \quad (109)$$

Then, we ensure at least \tilde{n} good columns exist. By Assumption 5.4, $(\alpha n)p_\star^k \rightarrow \infty$, the lower bound Eqn. (107) implies $\mu_n \rightarrow \infty$, and therefore $\exp(-\mu_n/8) \rightarrow 0$. In particular, for all large n we have $\mu_n/2 \geq \tilde{n}$ and

$$\mathbb{P}\left(N_{\text{good}} < \tilde{n} \mid \mathcal{G}\right) \leq \mathbb{P}\left(N_{\text{good}} \leq \frac{1}{2}\mu_n \mid \mathcal{G}\right) \leq \exp(-\mu_n/8). \quad (110)$$

Taking expectations and using $\mathbb{P}(E_n) \rightarrow 1$ yields

$$\mathbb{P}(N_{\text{good}} \geq \tilde{n}) \rightarrow 1. \quad (111)$$

On this event, choose any $J \subset J_n$ of size \tilde{n} consisting of good columns; then by definition of G_j we have $M_{ij} = 1$ for all $i \in I$ and $j \in J$, which proves the lemma. \square

E.3. Proof of Theorem 5.7

Proof of Theorem 5.7. Fix $\varepsilon > 0$. By the classical universal approximation theorem on \mathbb{R}^k (Cybenko, 1989), recalled as Theorem 5.1 in our paper, there exists an integer $\tilde{n}_\varepsilon < \infty$ and parameters $\{(a_r, \theta_r, b_r)\}_{r=1}^{\tilde{n}_\varepsilon}$ such that

$$\sup_{u \in K} \left| f(u) - \sum_{r=1}^{\tilde{n}_\varepsilon} a_r \sigma(\theta_r^\top u + b_r) \right| < \varepsilon. \quad (112)$$

Let I be the active input set from Assumption 5.3 and define $\tilde{f}(x) = f(x_I)$ and $\tilde{K} = \{x \in \mathbb{R}^d : x_I \in K, \|x_{I^c}\|_\infty \leq B\}$ as in the theorem. By Lemma 5.6, with probability $1 - o(1)$ there exists a set $J = \{j_1, \dots, j_{\tilde{n}_\varepsilon}\} \subset [n]$ such that $M_{ij_r} = 1$ for all $i \in I$ and all $r \leq \tilde{n}_\varepsilon$. On this event, we construct $F_n \in \mathcal{F}_n(M)$ by choosing parameters (a, b, θ) in Eqn. (7) as follows:

- For each $r \leq \tilde{n}_\varepsilon$, set $a_{j_r} := a_r$ and $b_{j_r} := b_r$.
- For each $r \leq \tilde{n}_\varepsilon$ and each $i \in I$, set $(\theta_{ij_r})_{i \in I} := \theta_r$, so that $\sum_{i \in I} \theta_{ij_r} x_i = \theta_r^\top x_I$.
- Set all remaining parameters a_j, b_j, θ_{ij} (for $j \notin J$ or $i \notin I$) equal to 0.

Because $M_{ij_r} = 1$ for all $i \in I$, these weights are not masked out, and thus for all $x \in \mathbb{R}^d$, we have

$$F_n(x) = \sum_{r=1}^{\tilde{n}_\varepsilon} a_r \sigma\left(b_r + \sum_{i \in I} \theta_{ij_r} x_i\right) = \sum_{r=1}^{\tilde{n}_\varepsilon} a_r \sigma(b_r + \theta_r^\top x_I). \quad (113)$$

Hence, for any $x \in \tilde{K}$ we have $x_I \in K$ and therefore by Eqn. (112),

$$|F_n(x) - \tilde{f}(x)| = \left| \sum_{r=1}^{\tilde{n}_\varepsilon} a_r \sigma(\theta_r^\top x_I + b_r) - f(x_I) \right| \leq \sup_{u \in K} \left| \sum_{r=1}^{\tilde{n}_\varepsilon} a_r \sigma(\theta_r^\top u + b_r) - f(u) \right| < \varepsilon. \quad (114)$$

Taking the supremum over $x \in \tilde{K}$ gives $\sup_{x \in \tilde{K}} |F_n(x) - \tilde{f}(x)| < \varepsilon$ on the same high-probability event. This proves

$$\mathbb{P}\left(\exists F_n \in \mathcal{F}_n(M) \text{ s.t. } \sup_{x \in \tilde{K}} |F_n(x) - \tilde{f}(x)| < \varepsilon\right) \rightarrow 1. \quad (115)$$

\square

F. Proof of Corollary 6.1

In this appendix, we prove Corollary 6.1. The argument works as follows: in the infinite-width Graphon-NTK regime, training the graphon-masked network induces the same output predictor as kernel (S)GD with the Graphon-NTK Gram matrix $K_{\mathcal{W}}$ on the fixed sample S ; we can then invoke Cao & Gu (2019, Cor. 3.10) with $K_{\mathcal{W}}$ in place of the finite-width NTK Gram matrix.

F.1. Preliminaries: predictions, Jacobian, and the empirical NTK Gram matrix

Let $S = \{(x_i, y_i)\}_{i=1}^m$ be i.i.d. samples from \mathcal{D} . Let $f_\theta : \mathcal{X} \rightarrow \mathbb{R}$ denote the masked network predictor. Define the prediction vector on the sample by

$$f(\theta) := (f_\theta(x_1), \dots, f_\theta(x_m))^\top \in \mathbb{R}^m. \quad (116)$$

Let the empirical risk be

$$\mathcal{L}_S(\theta) := \frac{1}{m} \sum_{i=1}^m \ell(f_\theta(x_i), y_i), \quad (117)$$

where ℓ is a differentiable surrogate loss like logistic or cross-entropy. Let $J(\theta) \in \mathbb{R}^{m \times p}$ denote the Jacobian of $f(\theta)$ with respect to parameters $\theta \in \mathbb{R}^p$: $J(\theta)_{i,:} = \nabla_\theta f_\theta(x_i)^\top$. Define the empirical NTK Gram matrix as

$$K(\theta) := J(\theta)J(\theta)^\top \in \mathbb{R}^{m \times m}. \quad (118)$$

Training dynamics. Under gradient flow $\dot{\theta}(t) = -\nabla_\theta \mathcal{L}_S(\theta(t))$, the chain rule gives the exact prediction dynamics

$$\dot{f}(t) = -K(\theta(t)) \nabla_f \mathcal{L}_S(f(t)), \quad (119)$$

where $\nabla_f \mathcal{L}_S(f) \in \mathbb{R}^m$ denotes the gradient of \mathcal{L}_S w.r.t. the prediction vector. For full-batch gradient descent $\theta_{t+1} = \theta_t - \eta \nabla_\theta \mathcal{L}_S(\theta_t)$, one obtains the standard first-order approximation in prediction space

$$f_{t+1} \approx f_t - \eta K(\theta_t) \nabla_f \mathcal{L}_S(f_t), \quad (120)$$

up to a second-order term controlled by smoothness of $f(\theta)$ and the stepsize η .

F.2. Graphon-NTK and lazy training regime

Let M_n be the PaI mask at width n , and suppose the associated step kernel converges to a limiting graphon $W_{M_n} \rightarrow \mathcal{W}$ (Section 4). Let $\Theta_{\mathcal{W}}$ denote the Graphon-NTK kernel defined in Pham et al. (2025), and define the deterministic Graphon-NTK Gram matrix on S by

$$(K_{\mathcal{W}})_{ij} := \Theta_{\mathcal{W}}(x_i, x_j). \quad (121)$$

We work under the following standard NTK lazy training assumption specialised to the graphon-masked setting.

Assumption F.1 (Infinite-width Graphon-NTK lazy training regime). Fix a training horizon $T \in \mathbb{N}$ for GD, or $T > 0$ for gradient flow and a stepsize schedule. As width $n \rightarrow \infty$, we assume:

1. **Deterministic kernel limit at initialization:** $K_n(\theta_{n,0}) \xrightarrow{\mathbb{P}} K_{\mathcal{W}}$ on the fixed sample S .
2. **Kernel stability along training (lazy training):** $\sup_{t \leq T} \|K_n(\theta_{n,t}) - K_n(\theta_{n,0})\|_{\text{op}} \xrightarrow{\mathbb{P}} 0$.
3. **Kernelised output predictor:** The network output predictor used in Cao & Gu (2019, Cor. 3.10) converges in law to the corresponding output predictor obtained by running kernel gradient descent using Gram matrix $K_{\mathcal{W}}$ on S with the same stepsize schedule.

Assumption F.1 places our analysis in the Graphon-NTK lazy training regime, in which the masked-network training dynamics on a fixed sample are governed by the limiting kernel $K_{\mathcal{W}}$. The assumption is supported by the same mechanism as in classical NTK analyses: under NTK scaling, each individual parameter moves by a vanishing amount as widths grow, which in turn keeps the network Jacobian, and hence the empirical NTK, nearly constant throughout training (Lee et al., 2019a; Arora et al., 2019b; Chizat et al., 2019).

To see the intuition, let $f(\theta) \in \mathbb{R}^m$ be the vector of predictions on the training set and let $J(\theta) \in \mathbb{R}^{m \times p}$ be the Jacobian with entries $J_{i,:} = \nabla_{\theta} f_{\theta}(x_i)^{\top}$. The empirical NTK Gram matrix is $K(\theta) = J(\theta)J(\theta)^{\top}$. Gradient flow satisfies the exact identity

$$\dot{f}(t) = -K(\theta(t)) \nabla_f \mathcal{L}_S(f(t)), \quad (122)$$

so kernel-like dynamics follow once $K(\theta(t))$ remains close to $K(\theta(0))$.

In sparse networks, weights are still independent at initialisation but have variances modulated by the layer graphons. Under bounded graphons and standard Lindeberg/LLN conditions, as used to establish the deterministic Graphon-NTK limit (Pham et al., 2025), forward activations remain $O(1)$ and backpropagated sensitivities obey the usual NTK scaling. Consequently, the gradient of the loss with respect to any single weight is of order $O((n_{\ell-1}n_{\ell})^{-1/2})$ in layer ℓ , so each parameter update is $o(1)$ as width grows. Aggregating across finitely many training steps, this implies that weights remain close to their initial values in a row-wise sense, pre-activations shift by $o(1)$, and therefore σ' remain stable. Since the Jacobian $J(\theta)$ is composed of products of these stable forward/backward factors, it follows that

$$\sup_{t \leq T} \|J(\theta(t)) - J(\theta(0))\|_{\text{op}} \rightarrow 0 \quad \text{and} \quad \sup_{t \leq T} \|K(\theta(t)) - K(\theta(0))\|_{\text{op}} \rightarrow 0, \quad (123)$$

for any fixed horizon T , yielding lazy training. A complete proof in the graphon-modulated setting would follow the standard NTK blueprint, with two additional technical steps: (i) concentration/CLT bounds for non-i.i.d. sums induced by the graphon variance profile, and (ii) uniform control over neurons and layers of forward and backward moments to ensure no small subset of nodes dominates the dynamics. Both are compatible with the bounded-graphon assumptions used to define the Graphon-NTK limit (Pham et al., 2025).

F.3. Proof of Corollary 6.1

We first recall from Section 3 of (Cao & Gu, 2019) the definition of the associated output predictor \hat{f} in Corollary 6.1. For a neural network with input $x \in \mathbb{R}^d$ and output $f_{\theta}(x)$ The associated output predictor \hat{f} is defined as

$$\hat{f}(x) := \mathbf{1}\{y \cdot f_{\theta}(x) < 0\}. \quad (124)$$

Proof of Corollary 6.1. Fix $\delta \in (0, e^{-1})$ and assume $\lambda_{\min}(K_{\mathcal{W}}) \geq \lambda_0 > 0$. By Assumption F.1, in the infinite-width Graphon-NTK (lazy) regime, the output predictor \hat{f} produced by training the graphon-masked network coincides with the output predictor produced by kernel gradient descent using Gram matrix $K_{\mathcal{W}}$ on the same sample S , with the same output convention as in Cao & Gu (2019, Cor. 3.10).

Therefore, we apply Cao & Gu (2019, Cor. 3.10) with the Gram matrix $\Theta^{(L)}$ in their bound replaced by $K_{\mathcal{W}}$. This yields that, with probability at least $1 - \delta$ over S ,

$$\mathbb{E}[\mathcal{L}_{0.1}^{\mathcal{D}}(\hat{f})] \leq CL \cdot \sqrt{\frac{y^{\top} K_{\mathcal{W}}^{-1} y}{m}} + C \sqrt{\frac{\log(1/\delta)}{m}},$$

where $y = (y_1, \dots, y_m)^{\top}$, $C > 0$ is the constant from Cao & Gu (2019), and L is the depth factor appearing there. This is exactly the bound in Eqn. (9). \square

G. Path-density proxy for the Graphon-NTK

From Section 4, masks generated by PaI method converge to a specific graphon, which describes the infinite-width limit of the induced sparsity. Given such converged graphon \mathcal{W} , (Pham et al., 2025) defines the associated *Graphon-NTK kernel* $\Theta_{\mathcal{W}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ as the deterministic infinite-width limit of the neural tangent kernel for the graphon-modulated network. On the sample S , and L -hidden-layer network, we define the *Graphon-NTK Gram matrix* $(K_{\mathcal{W}})_{ij} := \Theta_{\mathcal{W}}(x_i, x_j) \in \mathbb{R}^{m \times m}$:

$$K_{\mathcal{W}}(x, x') = \sum_{\ell=1}^L \int_0^1 \Sigma^{(\ell-1)}(u_{\ell-1}; x, x') du_{\ell-1} \int_0^1 \dot{\Sigma}^{(\ell)}(u_{\ell}; x, x') \left(\int_{[0,1]^{L-\ell+1}} \prod_{k=\ell+1}^{L+1} \mathcal{W}^{(k)}(u_{k-1}, u_k) \dot{\Sigma}^{(k)}(u_k; x, x') du_{\ell+1:L+1} \right) du_{\ell}, \quad (125)$$

G.1. Setup and notation

We consider a 2-hidden layer graphon network, with graphons $\mathcal{W}^{(1)}, \mathcal{W}^{(2)}, \mathcal{W}^{(3)} : [0, 1]^2 \rightarrow [0, 1]$. Let's u_0 index the input coordinate, u_1, u_2 index the two hidden layers, and u_3 index the output side. For inputs x, x' , the Graphon-NTK decomposes as $K_{\mathcal{W}} = K_{\mathcal{W}}^{(1)} + K_{\mathcal{W}}^{(2)}$ where

$$K_{\mathcal{W}}^{(1)}(x, x') = B_0(x, x') \int_0^1 \dot{\Sigma}^{(1)}(u_1; x, x') P_1(u_1; x, x') du_1, \quad (126)$$

$$K_{\mathcal{W}}^{(2)}(x, x') = B_1(x, x') \int_0^1 \dot{\Sigma}^{(2)}(u_2; x, x') P_2(u_2; x, x') du_2, \quad (127)$$

with

$$B_0(x, x') := \int_0^1 \Sigma^{(0)}(u_0; x, x') du_0, \quad B_1(x, x') := \int_0^1 \Sigma^{(1)}(u_1; x, x') du_1, \quad (128)$$

and the downstream passages

$$P_2(u_2; x, x') := \int_0^1 \mathcal{W}^{(3)}(u_2, u_3) \dot{\Sigma}^{(3)}(u_3; x, x') du_3, \quad (129)$$

$$P_1(u_1; x, x') := \int_0^1 \mathcal{W}^{(2)}(u_1, u_2) \dot{\Sigma}^{(2)}(u_2; x, x') P_2(u_2; x, x') du_2. \quad (130)$$

The pre-activation Gaussian cross-covariance at layer 1 is

$$\tilde{\Sigma}^{(1)}(u_1; x, x') = \int_0^1 \mathcal{W}^{(1)}(u_0, u_1) \Sigma^{(0)}(u_0; x, x') du_0. \quad (131)$$

Assumption G.1 (Activation regularity). σ is a.e. twice differentiable with $\|\sigma'\|_{\infty} \leq C_{\sigma,1}$ and $\|\sigma''\|_{\infty} \leq C_{\sigma,2}$.

The key steps of the proof are: (i) decompose $K_{\mathcal{W}} = K_{\mathcal{W}}^{(1)} + K_{\mathcal{W}}^{(2)}$ by layer contributions; (ii) bound activation covariances and derivative activation covariances by pre-activation covariances; (iii) recursively substitute graphon-weighted integrals and apply Fubini to isolate the path density term.

G.2. Two Gaussian comparison lemmas

Lemma G.2 (Gaussian covariance contraction). *Let (Z, Z') be centered jointly Gaussian. Then*

$$|\text{Cov}(\sigma(Z), \sigma(Z'))| \leq C_{\sigma,1}^2 |\text{Cov}(Z, Z')|. \quad (132)$$

Equivalently, if $F(s) := \mathbb{E}[\sigma(Z_s)\sigma(Z'_s)]$ where (Z_s, Z'_s) is centered Gaussian with fixed marginal variances and cross-covariance s , then

$$|F(s) - F(0)| \leq C_{\sigma,1}^2 |s|. \quad (133)$$

Proof. Fix the marginals and interpolate the cross-covariance: let (Z_t, Z'_t) be centered Gaussian with $\text{Cov}(Z_t, Z'_t) = t \text{Cov}(Z, Z')$ for $t \in [0, 1]$, and set $f(t) := \mathbb{E}[\sigma(Z_t)\sigma(Z'_t)]$. Then $f(1) - f(0) = \text{Cov}(\sigma(Z), \sigma(Z'))$ since $t = 0$ yields independence. A standard Gaussian differentiation identity gives $f'(t) = \text{Cov}(Z, Z') \mathbb{E}[\sigma'(Z_t)\sigma'(Z'_t)]$. Using $|\sigma'| \leq C_{\sigma,1}$ yields $|f'(t)| \leq C_{\sigma,1}^2 |\text{Cov}(Z, Z')|$, and integrating over $t \in [0, 1]$ gives the claim. \square

Lemma G.3 (Lipschitz dependence of $\dot{\Sigma}$ on cross-covariance). *Let (Z_s, Z'_s) be centered Gaussian with fixed marginal variances and cross-covariance s . Define $\dot{F}(s) := \mathbb{E}[\sigma'(Z_s)\sigma'(Z'_s)]$. Then \dot{F} is Lipschitz and*

$$|\dot{F}(s_1) - \dot{F}(s_2)| \leq C_{\sigma,2}^2 |s_1 - s_2|. \quad (134)$$

In particular, we have

$$|\dot{F}(s) - \dot{F}(0)| \leq C_{\sigma,2}^2 |s|. \quad (135)$$

Proof. We apply Lemma G.2 with the activation replaced by σ' . Indeed, $(\sigma')'$ exists a.e. and is bounded by $\|\sigma''\|_\infty \leq C_{\sigma,2}$, hence

$$|\text{Cov}(\sigma'(Z), \sigma'(Z'))| \leq C_{\sigma,2}^2 |\text{Cov}(Z, Z')|. \quad (136)$$

For the Gaussian family indexed by s , the centered covariance of $\sigma'(Z_s)$ and $\sigma'(Z'_s)$ equals $\dot{F}(s) - \dot{F}(0)$, which gives the desired inequality. \square

G.3. Path-density bounds for $K_{\mathcal{W}}^{(1)}$ and $K_{\mathcal{W}}^{(2)}$

Bounding $K_{\mathcal{W}}^{(1)}$ by path density. Since the activation derivative is bounded $\|\sigma'\|_\infty \leq C_{\sigma,1}$, then the expected correlation between activation derivative $\dot{\Sigma}^{(1)}(\sigma'(z), \sigma'(z')) \leq C_{\sigma,1}^2$. From Eqn. (129) and Eqn. (130), we have

$$0 \leq |P_1(u_1; x, x')| \leq C_{\sigma,1}^4 \int_{[0,1]^2} \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_2 du_3. \quad (137)$$

Next, let us fix u_1, x, x' . The quantity $\dot{\Sigma}^{(1)}(u_1; x, x')$ is of the form $\dot{F}(s)$ with $s = \tilde{\Sigma}^{(1)}(u_1; x, x')$. By Lemma G.3, we have

$$|\text{Cov}(\sigma'(z^{(1)}(u_1, x)), \sigma'(z^{(1)}(u_1, x')))| \leq C_{\sigma,2}^2 |\tilde{\Sigma}^{(1)}(u_1; x, x')|, \quad (138)$$

$$|\mathbb{E}[\sigma'(z^{(1)}(u_1, x))\sigma'(z^{(1)}(u_1, x'))] - \mathbb{E}[\sigma'(z^{(1)}(u_1, x))] \mathbb{E}[\sigma'(z^{(1)}(u_1, x'))]| \leq C_{\sigma,2}^2 |\tilde{\Sigma}^{(1)}(u_1; x, x')|, \quad (139)$$

$$|\dot{\Sigma}^{(1)}(u_1; x, x') - \mathbb{E}[\sigma'(z^{(1)}(u_1, x))] \mathbb{E}[\sigma'(z^{(1)}(u_1, x'))]| \leq C_{\sigma,2}^2 |\tilde{\Sigma}^{(1)}(u_1; x, x')|, \quad (140)$$

$$|\dot{\Sigma}^{(1)}(u_1; x, x')| \leq C_{\sigma,2}^2 |\tilde{\Sigma}^{(1)}(u_1; x, x')| + |\mathbb{E}[\sigma'(z^{(1)}(u_1, x))] \mathbb{E}[\sigma'(z^{(1)}(u_1, x'))]|, \quad (141)$$

Then, we have

$$|\dot{\Sigma}^{(1)}(u_1; x, x')| \leq C_{\sigma,2}^2 |\tilde{\Sigma}^{(1)}(u_1; x, x')| + c_1, \quad (142)$$

where c_1 is bounded by $C_{\sigma,1}^2$. Plugging Eqn. (142) and Eqn. (137) into Eqn. (126) gives

$$|K_{\mathcal{W}}^{(1)}(x, x')| \leq |B_0(x, x')| \int_0^1 \left(C_{\sigma,2}^2 |\tilde{\Sigma}^{(1)}(u_1; x, x')| + c_1 \right) P_1(u_1; x, x') du_1 \quad (143)$$

$$\leq |B_0(x, x')| C'_1 \left(\int_0^1 |\tilde{\Sigma}^{(1)}(u_1; x, x')| \int_{[0,1]^2} \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_2 du_3 du_1 \right), \quad (144)$$

where C'_1 is a constant depend on the $C_{\sigma,1}, C_{\sigma,2}, c_1$. Now, we substitute Eqn. (131) and use Fubini theorem:

$$\int_0^1 |\tilde{\Sigma}^{(1)}(u_1; x, x')| \int_{[0,1]^2} \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1 \quad (145)$$

$$\leq \int_0^1 \left(\int_0^1 \mathcal{W}^{(1)}(u_0, u_1) |\Sigma^{(0)}(u_0; x, x')| du_0 \right) \int_{[0,1]^2} \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1 \quad (146)$$

$$= \int_0^1 |\Sigma^{(0)}(u_0; x, x')| \left(\int_0^1 \mathcal{W}^{(1)}(u_0, u_1) \int_{[0,1]^2} \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1 \right) du_0 \quad (147)$$

$$= \int_0^1 |\Sigma^{(0)}(u_0; x, x')| \int_{[0,1]^3} \mathcal{W}^{(1)}(u_0, u_1) \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1 du_0. \quad (148)$$

Combining Eqn. (144) and Eqn. (148) yields

$$|K_{\mathcal{W}}^{(1)}(x, x')| \lesssim |B_0(x, x')| \left(\int_0^1 |\Sigma^{(0)}(u_0; x, x')| \times \int_{[0,1]^3} \mathcal{W}^{(1)}(u_0, u_1) \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1 du_0. \right) \quad (149)$$

Bounding $K_{\mathcal{W}}^{(2)}$ by path density. Similarly, using the bounded second derivative of activation function, we have

$$0 \leq |P_2(u_2; x, x')| \leq C_{\sigma,1}^2 \int_0^1 W^{(3)}(u_2, u_3) du_3, \quad (150)$$

Next, we fix u_2, x, x' . The quantity $\dot{\Sigma}^{(2)}(u_2; x, x')$ is of the form $\dot{F}(s)$ with $s = \tilde{\Sigma}^{(2)}(u_2; x, x')$. By Lemma G.3,

$$|\dot{\Sigma}^{(2)}(u_2; x, x')| \leq C_{\sigma,2}^2 |\tilde{\Sigma}^{(2)}(u_2; x, x')| + c_2. \quad (151)$$

Plugging Eqn. (150) and Eqn. (151) into Eqn. (127) gives

$$|K_{\mathcal{W}}^{(2)}(x, x')| \leq |B_1(x, x')| \int_0^1 \left(C_{\sigma,2}^2 |\tilde{\Sigma}^{(2)}(u_2; x, x')| + c_2 \right) P_2(u_2; x, x') du_2 \quad (152)$$

$$\leq |B_1(x, x')| C_2' \left(\int_0^1 |\tilde{\Sigma}^{(2)}(u_2; x, x')| \int_0^1 \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 \right), \quad (153)$$

where C_2' is a constant depend on the $C_{\sigma,1}, C_{\sigma,2}, c_2$. The pre-activation Gaussian cross-covariance at layer 2 is

$$\tilde{\Sigma}^{(2)}(u_2; x, x') = \int_0^1 \mathcal{W}^{(2)}(u_1, u_2) \Sigma^{(1)}(u_1; x, x') du_1. \quad (154)$$

With fixed u_1, x, x' , the quantity $\Sigma^{(1)}(u_1; x, x')$ is of the form $F(s)$ with $s = \tilde{\Sigma}^{(1)}(u_1; x, x')$. By Lemma G.2,

$$|\Sigma^{(1)}(u_1; x, x')| \leq C_{\sigma,1}^2 |\tilde{\Sigma}^{(1)}(u_1; x, x')| + c_1'. \quad (155)$$

where $c_1' = |\mathbb{E}[z^{(1)}(u_1; x)] \mathbb{E}[z^{(1)}(u_1; x')]|$. Combining Eqn. (155), Eqn. (154), and Eqn. (131) gives

$$|\tilde{\Sigma}^{(2)}(u_2; x, x')| \lesssim \int_{[0,1]^2} \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(1)}(u_0, u_1) |\Sigma^{(0)}(u_0; x, x')|, du_1 du_0. \quad (156)$$

Plugging Eqn. (156) into Eqn. (153) gives

$$|K_{\mathcal{W}}^{(2)}(x, x')| \lesssim |B_1(x, x')| \left(\int_0^1 |\Sigma^{(0)}(u_0; x, x')| \times \int_{[0,1]^3} \mathcal{W}^{(1)}(u_0, u_1) \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1 du_0. \right) \quad (157)$$

Combine the bounds for $K_{\mathcal{W}}^{(1)}$ and $K_{\mathcal{W}}^{(2)}$, and absorb all scaling terms we have

$$|K_{\mathcal{W}}(x, x')| \leq |K_{\mathcal{W}}^{(1)}(x, x')| + |K_{\mathcal{W}}^{(2)}(x, x')| \quad (158)$$

$$\lesssim \left(C_1 \left| \int_0^1 \Sigma^{(1)}(u_1, x, x') du_1 \right| + C_0 \left| \int_0^1 \Sigma^{(0)}(u_0, x, x') du_0 \right| \right) \times \quad (159)$$

$$\left(\int_0^1 |\Sigma^{(0)}(u_0; x, x')| \int_{[0,1]^3} \mathcal{W}^{(1)}(u_0, u_1) \mathcal{W}^{(2)}(u_1, u_2) \mathcal{W}^{(3)}(u_2, u_3) du_3 du_2 du_1 du_0. \right) \quad (160)$$

where C_0, C_1 are constant depending on $C_{\sigma,1}, C_{\sigma,2}$ and variance of activations and derivative activations $z^{(l)}$.