From Graphons to Graphlets: A Graph-Limit Lens on Sparse Neural Connectivity and Training

A central challenge for *network science for AI* is to understand how connectivity structure shapes learning in sparse neural systems. Sparse subnetworks can have the same parameter count yet train very differently, so task performance alone does not reveal which structural properties of connectivity matter for optimization. We argue that training dynamics should therefore be studied through a graph-theoretic lens: a pruning mask is not just a compression pattern, but a bipartite graph whose topology governs path structure, heterogeneity, and the routes through which information can flow at scale. The key question is then: *what is the correct graph-limit descriptor of neural connectivity as width grows?*

We study width-growing layer-wise masks $M_n \in \{0,1\}^{m_n \times r_n}$ as sequences of bipartite graphs and show that the answer is regime-dependent. In fixed-density settings, the associated step kernels admit a natural bipartite graphon description in the cut-distance sense [1, 2]. In shrinking-density settings, however, the raw adjacency picture collapses and ceases to resolve structure. This distinction matters for AI: if the asymptotic descriptor is chosen incorrectly, structurally meaningful sparse patterns can be washed out even when they still influence trainability.

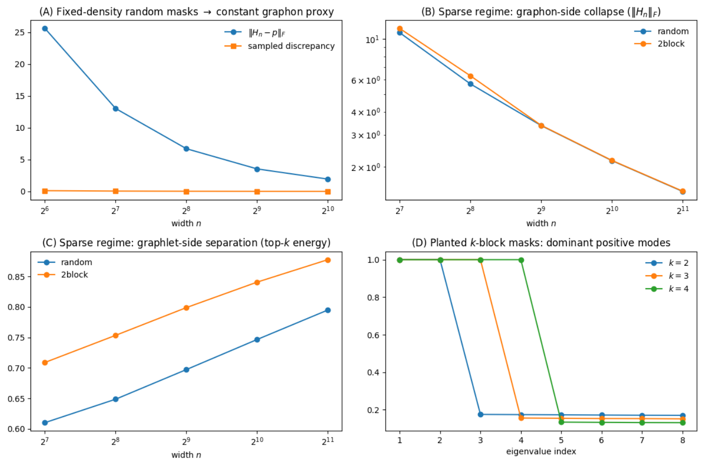

Figure 1: Regime map for width-growing neural masks. **A:** In fixed-density settings, random masks approach a constant bipartite graphon. **B:** Under vanishing density, graphon-side histogram norms decay for both random and planted sparse masks. **C:** Degree-normalized spectral summaries of the symmetrized bipartite mask graph still separate planted sparse masks from random sparse masks after graphon-side collapse. **D:** Planted k-block masks exhibit approximately k dominant positive normalized modes.

This perspective builds on recent progress connecting graph limits to pruning. Pham et al. proposed the *Graphon Limit Hypothesis* and a Graphon NTK for pruning masks in fixed-density regimes [4]. Our contribution is to sharpen that line into a *graph-limit regime map*: we identify when graphon language is appropriate, when it provably loses resolution, and why sparse regimes call for normalized operator-level descriptors in the spirit of Chung's graphlets framework [3].

On the theory side, we prove three structural results. First, random Bernoulli masks converge in probability to a constant bipartite graphon, giving a null model for unstructured fixed-density sparsity. Second, a dense discrepancy condition also forces convergence to the same constant limit, identifying a broad quasirandom regime in which mesoscale heterogeneity disappears asymptotically. Third, if the retained density $p_n \to 0$, then the unscaled mask kernel converges to zero in bipartite cut distance. Conceptually, this last theorem is the key point: once density vanishes, adjacency-level graphon summaries become asymptotically blind even though normalized structure may persist.

We support this picture with two empirical studies. First, the synthetic experiments in Fig. 1 mirror the theory: graphon-side observables are informative in fixed-density settings but decay under vanishing density, whereas degree-normalized spectral summaries continue to distinguish random sparse masks from planted modular sparse masks after graphon-side collapse. Second, on 48 pruned MNIST subnetworks (4 pruning methods $\times$ 4 sparsity levels $\times$ 3 seeds), we compare matched graphon-inspired and normalized spectral mask representations while controlling for pruning method and sparsity. Normalized spectral features are especially informative for medium-horizon optimization behavior: cross-validated R^2 for predicting 200-step relative loss drop improves from 0.67 to 0.81, and for predicting 200-step loss AUC from 0.54 to 0.76. Thus, the proposed regime split is not only asymptotic: it already helps explain finite-width differences in how sparse subnetworks train.

Overall, we propose a regime-aware graph-theoretic toolkit for sparse AI. Graphons provide the right first-order continuum objects for fixed-density neural connectivity, but after density collapse the structurally meaningful descriptors are normalized operator-level summaries that better capture modular sparse organization and its relation to training dynamics. More broadly, the work positions graph limits as a principled language for understanding when structure in AI systems is visible, when it is erased by the wrong asymptotic lens, and how network science can help turn sparse connectivity from an engineering artifact into an analyzable driver of learning.

References

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

[2] C. Borgs, J. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi. Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. *Adv. Math.*, 219(6):1801–1851, 2008.

[3] F. Chung. From quasirandom graphs to graph limits and graphlets. *Adv. Appl. Math.*, 56:135–174, 2014.

[4] H. Pham, T.-A. Ta, T. Jacobs, R. Burkholz, and L. Tran-Thanh. The Graphon Limit Hypothesis: Understanding Neural Network Pruning via Infinite Width Analysis. In *NeurIPS*, 2025.