

404

405 Appendix

406

407

Table of Contents

408	A Theoretical results	16
409	A.1 A hard example for GCN	16
410	A.2 A solution for GAT and CAT	17
411	B Synthetic experiments	22
412	B.1 Other experiments	24
413	C Dataset description	25
414	D Real data experiments	26
415	D.1 Experimental details	26
416	D.2 Additional results	27
417	E Open Graph Benchmark experiments	27
418	E.1 Experimental details	27
419	E.2 Additional results	28
420		
421		
422		

423 A Theoretical results

424 A.1 A hard example for GCN

425 In this subsection, we present a dataset and classification task for which GCN performs poorly. Note
426 that we follow the similar techniques and notation as [14], as described in the main paper.

427 We recall our data model. Fix $n, d \in \mathbb{N}$ and let $\varepsilon_1, \dots, \varepsilon_n$ be i.i.d uniformly sampled from $\{-1, 0, 1\}$.
428 Let $C_k = \{j \in [n] \mid \varepsilon_j = k\}$ for $k \in \{-1, 0, 1\}$. For each index $i \in [n]$, we set the feature vector
429 $\mathbf{X}_i \in \mathbb{R}^d$ as $\mathbf{X}_i \sim \mathcal{N}(\varepsilon_i \cdot \boldsymbol{\mu}, \mathbf{I} \cdot \sigma^2)$, where $\boldsymbol{\mu} \in \mathbb{R}^d$, $\sigma \in \mathbb{R}$ and $\mathbf{I} \in \{0, 1\}^{d \times d}$ is the identity matrix.
430 For a given pair $p, q \in [0, 1]$ we consider the stochastic adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$ defined
431 as follows. For $i, j \in [n]$ in the same class, we set $a_{ij} \sim \text{Ber}(p)$, and if i, j are in different classes,
432 we set $a_{ij} \sim \text{Ber}(q)$. We let $\mathbf{D} \in \mathbb{R}^{n \times n}$ be a diagonal matrix containing the degrees of the vertices.
433 We denote by $(\mathbf{X}, \mathbf{A}) \sim \text{CSBM}(n, p, q, \boldsymbol{\mu}, \sigma^2)$ a sample obtained according to the above random
434 process.

435 The task we wish to solve is classifying C_0 vs $C_{-1} \cup C_1$. Namely, we want our model φ to satisfy
436 $\varphi(\mathbf{X}_i) < 0$ if and only if $i \in C_0$. Moreover, note that the posed problem *is not linearly classifiable*.

437 To this end, we start by stating an assumption on the choice of parameters. This assumption is
438 necessary to achieve degree concentration in the graph.

439 **Assumption 1.** $p, q = \Omega(\log^2 n/n)$.

440 We now show the distribution of the convolved features. The following lemma can be easily obtained
441 using the techniques in [3].

442 **Lemma 3.** Fix p, q satisfying [Assumption 1](#). With probability at least $1 - o(1)$ over \mathbf{A} and $\{\varepsilon_i\}_i$,

$$(\mathbf{D}^{-1} \mathbf{A} \mathbf{X})_i \sim \mathcal{N}\left(\varepsilon_i \cdot \frac{p-q}{p+2q} \boldsymbol{\mu}, \frac{\sigma^2}{n(p+2q)}\right), \quad \forall i \in [n].$$

443 To prove the above lemma, we need the following definition of our high probability event.

444 **Definition 1.** We define the event \mathcal{E} as the intersection of the following events over \mathbf{A} and $\{\varepsilon_i\}_i$:

- 445 1. \mathcal{E}_1 is the event that $|C_0| = \frac{n}{3} \pm O(\sqrt{n \log n})$, $|C_1| = \frac{n}{3} \pm O(\sqrt{n \log n})$
446 and $|C_{-1}| = \frac{n}{3} \pm O(\sqrt{n \log n})$.
- 447 2. \mathcal{E}_2 is the event that for each $i \in [n]$, $\mathbf{D}_{ii} = \frac{n(p+2q)}{3} \left(1 \pm \frac{10}{\sqrt{\log n}}\right)$.
- 448 3. \mathcal{E}_3 is the event that for each $i \in [n]$ and $k \in \{-1, 0, 1\}$,

$$|N_i \cap C_k| = \begin{cases} \mathbf{D}_{ii} \cdot \frac{p}{p+2q} \cdot \left(1 \pm \frac{10}{\sqrt{\log n}}\right) & \text{if } i \in C_k \\ \mathbf{D}_{ii} \cdot \frac{q}{p+2q} \cdot \left(1 \pm \frac{10}{\sqrt{\log n}}\right) & \text{if } i \notin C_k \end{cases}.$$

449 The following lemma is a direct application of Chernoff bound and a union bound.

450 **Lemma 4.** With probability at least $1 - 1/\text{poly}(n)$ the event \mathcal{E} holds.

451 **Proof of Lemma 3.** By applying [Lemma 4](#), and conditioned on \mathcal{E} , for any $i \in [n]$

$$(\mathbf{D}^{-1} \mathbf{A} \mathbf{X})_i = \frac{1}{\mathbf{D}_{ii}} \sum_{j \in N_i} \mathbf{X}_j = \frac{1}{\mathbf{D}_{ii}} \left(\sum_{j \in N_i \cap C_{-1}} \mathbf{X}_j + \sum_{j \in N_i \cap C_0} \mathbf{X}_j + \sum_{j \in N_i \cap C_1} \mathbf{X}_j \right).$$

452 Using the definition of \mathcal{E} and properties of Gaussian distributions the lemma follows. \square

453 [Lemma 3](#) shows that essentially, the convolution reduced the variance and moved the means closer,
454 but the structure of the problem stayed exactly the same. Therefore, one layer of GCN cannot separate
455 C_0 from $C_{-1} \cup C_1$ with high probability.

456 **A.2 A solution for GAT and CAT**

457 In what follows, we show that GAT is able to handle the above classification task easily when the
 458 distance between the means is large enough. Then, we show how the additional convolution on the
 459 inputs to the score function improves the regime of perfect classification when the graph is not too
 460 noisy. Our main technical lemma considers a specific attention architecture and characterize the
 461 attention scores for our data model.

462 **Lemma 5.** Suppose that p, q satisfy [Assumption 1](#), $\|\boldsymbol{\mu}\| \geq \omega(\sigma\sqrt{\log n})$, fix the LeakyRelu constant
 463 $\beta \in (0, 1)$ and $R \in \mathbb{R}$. Then, there exists a choice of attention architecture Ψ such that with
 464 probability at least $1 - o_n(1)$ over the data $(\mathbf{X}, \mathbf{A}) \sim \text{CSBM}(n, p, q, \boldsymbol{\mu}, \sigma^2)$ the following holds.

$$\Psi(\mathbf{X}_i, \mathbf{X}_j) = \begin{cases} 10R\beta\|\boldsymbol{\mu}\|(1 \pm o(1)) & \text{if } i, j \in C_1^2 \\ -2R\|\boldsymbol{\mu}\|(1 + 2\beta)(1 \pm o(1)) & \text{if } i, j \in C_{-1}^2 \\ -2R\|\boldsymbol{\mu}\|(1 + 5\beta)(1 \pm o(1)) & \text{if } i \in C_1, j \in C_{-1} \\ 10R\beta\|\boldsymbol{\mu}\|(1 \pm o(1)) & \text{if } i \in C_{-1}, j \in C_1 \\ -\frac{R}{2}\|\boldsymbol{\mu}\|(1 - 21\beta)(1 \pm o(1)) & \text{if } i \in C_0, j \in C_1 \\ -\frac{R}{2}\|\boldsymbol{\mu}\|(1 - 11\beta)(1 \pm o(1)) & \text{if } i \in C_0, j \in C_{-1} \\ -\frac{R}{2}\|\boldsymbol{\mu}\|(1 - 5\beta)(1 \pm o(1)) & \text{if } i \in C_1, j \in C_0 \\ -\frac{R}{2}\|\boldsymbol{\mu}\|(1 - 5\beta)(1 \pm o(1)) & \text{if } i \in C_{-1}, j \in C_0 \\ 2R\beta\|\boldsymbol{\mu}\|(1 \pm o(1)) & \text{if } i, j \in C_0^2 \end{cases}.$$

465 **Proof.** We consider as an ansatz the following two layer architecture Ψ .

$$\tilde{\mathbf{w}} \stackrel{\text{def}}{=} \frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|}, \quad \mathbf{S} \stackrel{\text{def}}{=} \begin{bmatrix} 1 & 1 \\ -1 & -1 \\ 1 & -1 \\ -1 & 1 \\ 0 & 1 \\ 1 & 0 \\ 0 & -1 \\ -1 & 0 \end{bmatrix}, \quad \mathbf{b} \stackrel{\text{def}}{=} \begin{bmatrix} -3/2 \\ -3/2 \\ -3/2 \\ -3/2 \\ -1/2 \\ -1/2 \\ -1/2 \\ -1/2 \end{bmatrix} \cdot \|\boldsymbol{\mu}\|, \quad \mathbf{r} \stackrel{\text{def}}{=} R \cdot \begin{bmatrix} 2 \\ -2 \\ -2 \\ 2 \\ -1 \\ -1 \\ -1 \\ -1 \end{bmatrix},$$

466 where $R > 0$ is an arbitrary scaling parameter. The output of the attention model is defined as

$$\Psi(\mathbf{X}_i, \mathbf{X}_j) \stackrel{\text{def}}{=} \mathbf{r}^T \cdot \text{LeakyRelu} \left(\mathbf{S} \cdot \begin{bmatrix} \tilde{\mathbf{w}}^T \mathbf{X}_i \\ \tilde{\mathbf{w}}^T \mathbf{X}_j \end{bmatrix} + \mathbf{b} \right).$$

467 Let $\boldsymbol{\Delta}_{ij} \stackrel{\text{def}}{=} \mathbf{S} \cdot \begin{bmatrix} \tilde{\mathbf{w}}^T \mathbf{X}_i \\ \tilde{\mathbf{w}}^T \mathbf{X}_j \end{bmatrix} + \mathbf{b} \in \mathbb{R}^8$, and note that for each element $t \in [8]$ of $\boldsymbol{\Delta}_{ij}$, we have that

468 $(\boldsymbol{\Delta}_{ij})_t = \mathbf{S}_{t,1} \tilde{\mathbf{w}}^T \mathbf{X}_i + \mathbf{S}_{t,2} \tilde{\mathbf{w}}^T \mathbf{X}_j + \mathbf{b}_t$. Note that the random variable $(\boldsymbol{\Delta}_{ij})_t$ is distributed as
 469 follows:

$$(\boldsymbol{\Delta}_{ij})_t \sim \begin{cases} \mathcal{N}((\mathbf{S}_{t,1} + \mathbf{S}_{t,2})\tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i, j \in C_1^2 \\ \mathcal{N}(-(\mathbf{S}_{t,1} + \mathbf{S}_{t,2})\tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i, j \in C_{-1}^2 \\ \mathcal{N}((\mathbf{S}_{t,1} - \mathbf{S}_{t,2})\tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i \in C_1, j \in C_{-1} \\ \mathcal{N}(-(\mathbf{S}_{t,1} - \mathbf{S}_{t,2})\tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i \in C_{-1}, j \in C_1 \\ \mathcal{N}(\mathbf{S}_{t,2} \tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i \in C_0, j \in C_1 \\ \mathcal{N}(-\mathbf{S}_{t,2} \tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i \in C_0, j \in C_{-1} \\ \mathcal{N}(\mathbf{S}_{t,1} \tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i \in C_1, j \in C_0 \\ \mathcal{N}(-\mathbf{S}_{t,1} \tilde{\mathbf{w}}^T \boldsymbol{\mu} + \mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i \in C_{-1}, j \in C_0 \\ \mathcal{N}(\mathbf{b}_t, \|\mathbf{S}_{t,*}\|^2 \sigma^2) & \text{if } i, j \in C_0^2 \end{cases}.$$

470 Therefore, for a fixed $i, j \in [n]^2$ we have that the entries of $\boldsymbol{\Delta}_{ij}$ are distributed as follows (where we
 471 use \mathcal{N}_x^y as abbreviation for the Gaussian $\mathcal{N}(x, y)$)

$$\left[\mathcal{N}_{\frac{\|\boldsymbol{\mu}\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{-7\|\boldsymbol{\mu}\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{-3\|\boldsymbol{\mu}\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{-3\|\boldsymbol{\mu}\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{\|\boldsymbol{\mu}\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\boldsymbol{\mu}\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{-3\|\boldsymbol{\mu}\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{-3\|\boldsymbol{\mu}\|}{2}}^{\sigma^2} \right] \quad \text{for } i, j \in C_1^2,$$

$$\begin{aligned}
& \left[\mathcal{N}_{-\frac{7\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_{-1}^2, \\
& \left[\mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{7\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_1 \times C_{-1}, \\
& \left[\mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{7\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_{-1} \times C_1, \\
& \left[\mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_0 \times C_1, \\
& \left[\mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_0 \times C_{-1}, \\
& \left[\mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_1 \times C_0, \\
& \left[\mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{5\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{\frac{\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_{-1} \times C_0, \\
& \left[\mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{4\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{3\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \quad \mathcal{N}_{-\frac{\|\mu\|}{2}}^{\sigma^2} \right] && \text{for } i, j \in C_0^2,
\end{aligned}$$

472 Next, we will use the following lemma regarding LeakyRelu concentration.

473 **Lemma 6** (Lemma A.6 in [14]). Fix $s \in \mathbb{N}$, and let z_1, \dots, z_s be jointly Gaussian random variables
474 with marginals $z_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$. There exists an absolute constant $C > 0$ such that with probability
475 at least $1 - o_s(1)$, we have

$$\text{LeakyRelu}(z_i) = \text{LeakyRelu}(\mu_i) \pm C\sigma_i\sqrt{\log s}, \quad \text{for all } i \in [s].$$

476 Using Lemma 6 with the assumption on $\|\mu\|$ and a union bound, we have that with probability at
477 least $1 - o_n(1)$, $\text{LeakyRelu}(\Delta_{ij})$ is (up to $1 \pm o(1)$)

$$\begin{aligned}
& \left[\frac{\|\mu\|}{2} \quad -\frac{7\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \right] && \text{for } i, j \in C_1^2, \\
& \left[-\frac{7\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \right] && \text{for } i, j \in C_{-1}^2, \\
& \left[-\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{7\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \right] && \text{for } i, j \in C_1 \times C_{-1}, \\
& \left[-\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{7\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \right] && \text{for } i, j \in C_{-1} \times C_1, \\
& \left[-\frac{\beta\|\mu\|}{2} \quad -\frac{5\beta\|\mu\|}{2} \quad -\frac{5\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{\|\beta\mu\|}{2} \right] && \text{for } i, j \in C_0 \times C_1, \\
& \left[-\frac{5\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{5\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \right] && \text{for } i, j \in C_0 \times C_{-1}, \\
& \left[-\frac{\beta\|\mu\|}{2} \quad -\frac{5\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{5\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad \frac{\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \right] && \text{for } i, j \in C_1 \times C_0, \\
& \left[-\frac{5\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{5\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{\|\beta\mu\|}{2} \quad \frac{\|\mu\|}{2} \right] && \text{for } i, j \in C_{-1} \times C_0, \\
& \left[-\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{3\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \quad -\frac{\beta\|\mu\|}{2} \right] && \text{for } i, j \in C_0^2.
\end{aligned}$$

478 Then,

$$\mathbf{r}^T \cdot \text{LeakyRelu}(\Delta_{ij}) = \begin{cases} 10R\beta\|\mu\|(1 \pm o(1)) & \text{if } i, j \in C_1^2 \\ -2R\|\mu\|(1 + 2\beta)(1 \pm o(1)) & \text{if } i, j \in C_{-1}^2 \\ -2R\|\mu\|(1 + 5\beta)(1 \pm o(1)) & \text{if } i \in C_1, j \in C_{-1} \\ 10R\beta\|\mu\|(1 \pm o(1)) & \text{if } i \in C_{-1}, j \in C_1 \\ -\frac{R}{2}\|\mu\|(1 - 21\beta)(1 \pm o(1)) & \text{if } i \in C_0, j \in C_1 \\ -\frac{R}{2}\|\mu\|(1 - 11\beta)(1 \pm o(1)) & \text{if } i \in C_0, j \in C_{-1} \\ -\frac{R}{2}\|\mu\|(1 - 5\beta)(1 \pm o(1)) & \text{if } i \in C_1, j \in C_0 \\ -\frac{R}{2}\|\mu\|(1 - 5\beta)(1 \pm o(1)) & \text{if } i \in C_{-1}, j \in C_0 \\ 2R\beta\|\mu\|(1 \pm o(1)) & \text{if } i, j \in C_0^2 \end{cases},$$

479 and the proof is complete. \square

480 Next we will define our high probability event.

481 **Definition 2.** $\mathcal{E}' \stackrel{\text{def}}{=} \mathcal{E} \cap \mathcal{E}^*$, where \mathcal{E}^* is the event that for a fixed $\mathbf{w} \in \mathbb{R}^d$, all $i \in [n]$ satisfy
 482 $|\mathbf{w}^T \mathbf{X}_i - \mathbf{E}[\mathbf{w}^T \mathbf{X}_i]| \leq 10\sigma \|\mathbf{w}\|_2 \sqrt{\log n}$.

483 The following lemma is obtained by using [Lemma 4](#) with standard Gaussian concentration and a
 484 union bound.

485 **Lemma 7.** With probability at least $1 - 1/\text{poly}(n)$ event \mathcal{E}' holds.

486 **Corollary 8.** Suppose that p, q satisfy [Assumption 1](#), $\|\boldsymbol{\mu}\| = \omega(\sigma\sqrt{\log n})$ and fix $R \in \mathbb{R}$. Then,
 487 there exists a choice of attention architecture Ψ such that with probability $1 - o_n(1)$ over $(\mathbf{A}, \mathbf{X}) \sim$
 488 $\text{CSBM}(n, p, q, \boldsymbol{\mu}, \sigma^2)$ it holds that

$$\gamma_{ij} = \begin{cases} \frac{3}{np}(1 \pm o(1)) & \text{if } i, j \in C_0^2 \cup C_1^2 \\ \frac{3}{nq}(1 \pm o(1)) & \text{if } i, j \in C_{-1} \times C_1 \\ \frac{3}{nq} \exp(-\Theta(R\|\boldsymbol{\mu}\|)) & \text{if } i, j \in C_{-1} \times C_{-1} \cup C_0 \\ \frac{3}{np} \exp(-\Theta(R\|\boldsymbol{\mu}\|)) & \text{otherwise} \end{cases},$$

489 where R is a parameter of the architecture.

490 **Proof.** The proof is immediate. First applying the ansatz from [Lemma 5](#) with $\beta < 1/25$, [Lemma 7](#)
 491 and a union bound. Using the definition of γ_{ij} concludes the proof. \square

492 Next, we prove [Theorem 1](#) that the model distinguish nodes from C_0 for any choice of p, q satisfying
 493 [Assumption 1](#). We restate the theorem for convince.

494 **Theorem 9** (Formal restatement of [Theorem 1](#)). Suppose that p, q satisfy [Assumption 1](#) and $\|\boldsymbol{\mu}\|_2 =$
 495 $\omega(\sigma\sqrt{\log n})$. Then, there exists a choice of attention architecture Ψ such that with probability at least
 496 $1 - o_n(1)$ over the data $(\mathbf{X}, \mathbf{A}) \sim \text{CSBM}(n, p, q, \boldsymbol{\mu}, \sigma^2)$, the estimator

$$\hat{x}_i \stackrel{\text{def}}{=} \sum_{j \in N_i} \gamma_{ij} \tilde{\mathbf{w}}^T \mathbf{X}_j + b \text{ where } \tilde{\mathbf{w}} = \boldsymbol{\mu} / \|\boldsymbol{\mu}\|, b = -\|\boldsymbol{\mu}\|/2$$

497 satisfies $\hat{x}_i < 0$ if and only if $i \in C_0$.

498 **Proof.** Let Ψ be the architecture from [Corollary 8](#) and let R satisfy $R\|\boldsymbol{\mu}\|_2 = \omega(1)$. We will
 499 compute the mean and variance of the estimator \hat{x}_i conditioned on \mathcal{E}' . Suppose that $i \in C_0$. By using
 500 [Corollary 8](#), [Definition 2](#) and our assumption on $\|\boldsymbol{\mu}\|$ and R , we have

$$\max \left\{ \frac{3}{np} \exp(-\Theta(R\|\boldsymbol{\mu}\|)), \frac{3}{nq} \exp(-\Theta(R\|\boldsymbol{\mu}\|)) \right\} = o \left(\frac{1}{n(p+2q)} \right),$$

501 and therefore

$$\begin{aligned} \mathbf{E}[\hat{x}_i \mid \mathcal{E}'] &= \mathbf{E} \left[\sum_{k \in \{-1, 0, 1\}} \sum_{j \in N_i \cap C_k} \gamma_{ij} \tilde{\mathbf{w}}^T \mathbf{X}_j \mid \mathcal{E}' \right] - \frac{\|\boldsymbol{\mu}\|}{2} \\ &= \mathbf{E}[|C_0 \cap N_i| \mid \mathcal{E}'] \left(\pm \frac{3}{np} (1 \pm o(1)) \cdot 10\sigma \sqrt{\log n} \right) \\ &\quad + \mathbf{E}[|C_1 \cap N_i| \mid \mathcal{E}'] \left(o \left(\frac{1}{n(p+2q)} \right) \cdot (\|\boldsymbol{\mu}\| \pm 10\sigma \sqrt{\log n}) \right) \\ &\quad + \mathbf{E}[|C_{-1} \cap N_i| \mid \mathcal{E}'] \left(o \left(\frac{1}{n(p+2q)} \right) \cdot (-\|\boldsymbol{\mu}\| \pm 10\sigma \sqrt{\log n}) \right) - \frac{\|\boldsymbol{\mu}\|}{2} \\ &= -\frac{\|\boldsymbol{\mu}\|}{2} (1 \pm o(1)). \end{aligned}$$

502 By similar reasoning we have that for $i \in C_{-1} \cup C_1$, $\mathbf{E} [\hat{x}_i \mid \mathcal{E}'] = \frac{\|\boldsymbol{\mu}\|}{2}(1 \pm o(1))$.

503 Next, we claim that for each $i \in [n]$ the random variable \hat{x}_i given \mathcal{E}' is sub-Gaussian with a small
 504 sub-Gaussian constant compared to the above expectation. The following lemma is a straightforward
 505 adaptation of Lemma A.11 in [14], and we provide its proof for completeness.

506 **Lemma 10.** Conditioned on \mathcal{E}' , the random variables $\{\hat{x}_i\}_i$ are sub-Gaussian with parameter
 507 $\tilde{\sigma}_i^2 = O\left(\frac{\sigma^2}{np}\right)$ if $i \in C_0 \cup C_1$ and $\tilde{\sigma}_i^2 = O\left(\frac{\sigma^2}{nq}\right)$ otherwise.

508 **Proof.** Fix $i \in [n]$, and write $\mathbf{X}_i = \varepsilon_i \boldsymbol{\mu} + \sigma \mathbf{g}_i$ where $\mathbf{g}_i \sim \mathcal{N}(0, \mathbf{I}_d)$, and ε_i denotes the class
 509 membership. Consider \hat{x}_i as a function of $\mathbf{g} = [\mathbf{g}_1 \circ \mathbf{g}_2 \circ \dots \circ \mathbf{g}_n] \in \mathbb{R}^{nd}$, where \circ denotes vertical
 510 concatenation. Namely, consider the function

$$\hat{x}_i = f_i(\mathbf{g}) \stackrel{\text{def}}{=} \sum_{j \in N_i} \gamma_{ij}(\mathbf{g}) \tilde{\mathbf{w}}^T (\varepsilon_j \boldsymbol{\mu} + \sigma \mathbf{g}_j) - \|\boldsymbol{\mu}\|/2, \quad i \in [n].$$

511 Since $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{nd})$, proving that \hat{x}_i given \mathcal{E}' is sub-Gaussian for each $i \in [n]$, reduces to
 512 showing that the function $f_i : \mathbb{R}^{nd} \rightarrow \mathbb{R}$ is Lipschitz over $E \subseteq \mathbb{R}^{nd}$ defined by \mathcal{E}' and the relation
 513 $\mathbf{X}_i = \varepsilon_i \boldsymbol{\mu} + \sigma \mathbf{g}_i$. That is, $E \stackrel{\text{def}}{=} \{\mathbf{g} \in \mathbb{R}^{nd} \mid |\tilde{\mathbf{w}}^T \mathbf{g}_i| \leq 10\sqrt{\log n}, \forall i \in [n]\}$. Specifically, we show
 514 that conditioning on the event \mathcal{E}' (which restricts $\mathbf{g} \in E$), the Lipschitz constant L_{f_i} of f_i satisfies
 515 $L_{f_i} = O\left(\frac{\sigma}{\sqrt{np}}\right)$ for $i \in C_0 \cup C_1$ and $L_{f_i} = O\left(\frac{\sigma}{nq}\right)$ otherwise, and hence proving the claim.

516 To compute the Lipschitz constant of $f_i(\mathbf{g})$ for $i \in [n]$, let us denote $\mathbf{X} = [\mathbf{X}_1 \circ \mathbf{X}_2 \circ \dots \circ \mathbf{X}_n]$ and
 517 consider the function

$$\tilde{f}_i(\mathbf{X}) \stackrel{\text{def}}{=} \sum_{j \in N_i} \gamma_{ij}(\mathbf{X}) \tilde{\mathbf{w}}^T \mathbf{X}_j, \quad i \in [n]$$

Let us assume without loss of generality that $i \in C_0$ (the cases for $i \in C_1$ and $i \in C_{-1}$ are obtained
 identically). Conditioning on the event \mathcal{E}' , which imposes the restriction that $\mathbf{X} \in \tilde{E}$ where

$$\tilde{E} \stackrel{\text{def}}{=} \left\{ \mathbf{X} \in \mathbb{R}^{nd} \mid |\mathbf{X}_i - \varepsilon_i \boldsymbol{\mu}| \leq 10\sigma\sqrt{\log n}, \forall i \in [n] \right\}.$$

518 Conditioning on \mathcal{E}' (which restricts $\mathbf{X}, \mathbf{X}' \in \tilde{E}$), using [Corollary 8](#) and recalling that R satisfies
 519 $R\|\boldsymbol{\mu}\|_2 = \omega(1)$, we get⁵

$$\begin{aligned} & \left| \tilde{f}_i(\mathbf{X}) - \tilde{f}_i(\mathbf{X}') \right| \\ & \simeq \left| \sum_{j \in N_i \cap C_0} \frac{3}{np} \tilde{\mathbf{w}}^T (\mathbf{X}_j - \mathbf{X}'_j) + \sum_{j \in N_i \cap C_1} \frac{3}{np} \cdot e^{-\Theta(R\|\boldsymbol{\mu}\|_2)} \tilde{\mathbf{w}}^T (\mathbf{X}_j - \mathbf{X}'_j) + \sum_{j \in N_i \cap C_{-1}} \frac{3}{np} \cdot e^{-\Theta(R\|\boldsymbol{\mu}\|_2)} \tilde{\mathbf{w}}^T (\mathbf{X}_j - \mathbf{X}'_j) \right| \\ & = \left| \left[\begin{array}{l} \frac{3}{np} (1 \pm o(1)) \tilde{\mathbf{w}} \\ \frac{3}{np} \exp(-\Theta(R\|\boldsymbol{\mu}\|_2)) (1 \pm o(1)) \tilde{\mathbf{w}} \\ \frac{3}{np} \exp(-\Theta(R\|\boldsymbol{\mu}\|_2)) (1 \pm o(1)) \tilde{\mathbf{w}} \\ 0 \end{array} \right]_{j \in [n]}^T (\mathbf{X} - \mathbf{X}') \right| \\ & \leq \left\| \left[\begin{array}{l} \frac{3}{np} (1 \pm o(1)) \tilde{\mathbf{w}} \\ \frac{3}{np} \exp(-\Theta(R\|\boldsymbol{\mu}\|_2)) (1 \pm o(1)) \tilde{\mathbf{w}} \\ \frac{3}{np} \exp(-\Theta(R\|\boldsymbol{\mu}\|_2)) (1 \pm o(1)) \tilde{\mathbf{w}} \\ 0 \end{array} \right]_{j \in [n]} \right\|_2 \|\mathbf{X} - \mathbf{X}'\|_2 \\ & \leq \sqrt{\frac{3}{np}} (1 + o(1)) \|\tilde{\mathbf{w}}\|_2 \|\mathbf{X} - \mathbf{X}'\|_2 \\ & = \sqrt{\frac{3}{np}} (1 + o(1)) \|\mathbf{X} - \mathbf{X}'\|_2. \end{aligned}$$

⁵We drop the $(1 \pm o(1))$ in the first line of the computation for compactness and use \simeq as notation.

520 This shows the Lipschitz constant of $\tilde{f}_i(\mathbf{X})$ over \tilde{E} satisfies $L_{\tilde{f}_i} = O\left(\frac{1}{\sqrt{np}}\right)$. On the other hand, by
 521 viewing \mathbf{X} as a function of \mathbf{g} , it is straightforward to see that the function $h(\mathbf{g}) : \mathbb{R}^{nd} \rightarrow \mathbb{R}^{nd}$ defined
 522 by $h(\mathbf{g}) \stackrel{\text{def}}{=} \mathbf{X}(\mathbf{g})$ has Lipschitz constant $L_h = \sigma$, as

$$\|h(\mathbf{g}) - h(\mathbf{g}')\|_2 = \|\varepsilon\boldsymbol{\mu} + \sigma\mathbf{g} - (\varepsilon\boldsymbol{\mu} + \sigma\mathbf{g}')\|_2 = \sigma\|\mathbf{g} - \mathbf{g}'\|_2.$$

523 Therefore, since $f_i(\mathbf{g}) = \tilde{f}_i(h(\mathbf{g}))$ and $\mathbf{g} \in E$ if and only if $\mathbf{X} \in \tilde{E}$, we have that, conditioning on \mathcal{E}' ,
 524 the function $\hat{x}_i = f_i(\mathbf{g})$ is Lipschitz continuous with Lipschitz constant $L_{f_i} = L_{\tilde{f}_i}L_h = O\left(\frac{\sigma}{\sqrt{np}}\right)$.
 525 Since $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{nd})$, we know that \hat{x}_i is sub-Gaussian with sub-Gaussian constant $\tilde{\sigma}^2 = L_{f_i}^2 =$
 526 $O\left(\frac{\sigma^2}{np}\right)$. \square

527 The following lemma will be used for bounding the misclassification probability.

528 **Lemma 11** ([30]). Let x_1, \dots, x_n be sub-Gaussian random variables with the same mean and
 529 sub-Gaussian parameter $\tilde{\sigma}^2$. Then,

$$\mathbf{E} \left[\max_{i \in [n]} (x_i - \mathbf{E}[x_i]) \right] \leq \tilde{\sigma} \sqrt{2 \log n}.$$

530 Moreover, for any $t > 0$

$$\Pr \left[\max_{i \in [n]} (x_i - \mathbf{E}[x_i]) > t \right] \leq 2n \exp \left(-\frac{t^2}{2\tilde{\sigma}^2} \right).$$

531 We bound the probability of misclassification

$$\Pr \left[\max_{i \in C_0} \hat{x}_i \geq 0 \right] \leq \Pr \left[\max_{i \in C_0} \hat{x}_i > t + \mathbf{E}[\hat{x}_i] \right],$$

532 for $t < |\mathbf{E}[\hat{x}_i]| = \frac{\|\boldsymbol{\mu}\|_2}{2}(1 \pm o(1))$. By **Lemma 10**, picking $t = \Theta\left(\sigma\sqrt{\log|C_0|}\right)$ and applying
 533 **Lemma 11** implies that the above probability is $1/\text{poly}(n)$.

534 Similarly for class $C_1 \cup C_{-1}$ we have that the misclassification probability is

$$\begin{aligned} \Pr \left[\min_{i \in C_1 \cup C_{-1}} \hat{x}_i \leq 0 \right] &= \Pr \left[-\max_{i \in C_1 \cup C_{-1}} (-\hat{x}_i) \leq 0 \right] = \Pr \left[\max_{i \in C_1 \cup C_{-1}} (-\hat{x}_i) \geq 0 \right] \\ &\leq \Pr \left[\max_{i \in C_1 \cup C_{-1}} -\hat{x}_i > t - \mathbf{E}[\hat{x}_i] \right], \end{aligned}$$

535 for $t < \mathbf{E}[\hat{x}_i]$. Picking $t = \Theta\left(\sigma\sqrt{\log|C_1 \cup C_{-1}|}\right)$ and applying **Lemma 11** and a union bound over
 536 the misclassification probabilities of both classes conclude the proof of the corollary. \square

537 Combining **Theorem 9** with **Lemma 3**, we immediately get **Corollary 2** which we restate below.

538 **Corollary 12.** Suppose $p, q = \Omega(\log^2 n/n)$ and $\|\boldsymbol{\mu}\| \geq \omega\left(\sigma\sqrt{\frac{(p+2q)\log n}{n(p-q)^2}}\right)$. Then, there is a choice
 539 of attention architecture Ψ such that, with probability at least $1 - o(1)$ over the data $(\mathbf{X}, \mathbf{A}) \sim$
 540 $\text{CSBM}(n, p, q, \boldsymbol{\mu}, \sigma^2)$, CAT separates nodes C_0 from $C_1 \cup C_{-1}$.

541 B Synthetic experiments

542 In this section, we present the complete results for the synthetic data experiments of §6.1. First, we
 543 describe the parameterization we use for the 1-layer GCN, GAT, and CAT models; then, we verify the
 544 behavior of the normalized score function (γ_{ij}) matches that of the theory presented in Corollary 8.
 545 In particular, we visualize the average of the following three groups of gammas (Fig. 4):

- 546 • Gammas γ_{ij} included in $i, j \in C_0^2 \cup C_1^2$. Solid lines.
- 547 • Gammas γ_{ij} included in $i, j \in C_{-1} \times C_1$. Dashed lines.
- 548 • The rest of gammas. Dotted lines.

549 For completeness, we also include the empirical results that validate Theorem 1 and Corollary 2,
 550 which were discussed already in §6.1.

551 **Experimental setup.** We assume the following parametrization for the 1-layer GCN, GAT, and CAT:

$$\mathbf{h}'_i = \left(\sum_{j \in N_i} \gamma_{ij} \tilde{\mathbf{w}}^T \mathbf{X}_j \right) - C \cdot \|\boldsymbol{\mu}\|/2, \quad (8)$$

552 where N_i are the set of neighbors of node i , \mathbf{X}_j are the features of node j —obtained from the
 553 CSBM described in §4, and \mathbf{h}'_i are the logits of the prediction of node i . Note that for GCN we have
 554 $\gamma_{ij} = \frac{1}{|N_i^*|}$. Otherwise, we consider the following parameterization of the score function Ψ :

$$\gamma_{ij} = \frac{\exp(\Psi(\mathbf{h}_i, \mathbf{h}_j))}{\sum_{k \in N_i^*} \exp(\Psi(\mathbf{h}_i, \mathbf{h}_k))} \quad \text{where} \quad (9)$$

$$\Psi(\mathbf{h}_i, \mathbf{h}_j) \stackrel{\text{def}}{=} \mathbf{r}^T \cdot \text{LeakyRelu} \left(\mathbf{S} \cdot \begin{bmatrix} \tilde{\mathbf{w}}^T \mathbf{h}_i \\ \tilde{\mathbf{w}}^T \mathbf{h}_j \end{bmatrix} + \mathbf{b} \right). \quad (10)$$

555 For these experiments, we define the parameters $\tilde{\mathbf{w}}$, \mathbf{S} , \mathbf{b} and \mathbf{r} as in the proofs in Appendix A:

$$\tilde{\mathbf{w}} \stackrel{\text{def}}{=} \frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|}, \quad \mathbf{S} \stackrel{\text{def}}{=} \begin{bmatrix} 1 & 1 \\ -1 & -1 \\ 1 & -1 \\ -1 & 1 \\ 0 & 1 \\ 1 & 0 \\ 0 & -1 \\ -1 & 0 \end{bmatrix}, \quad \mathbf{b} \stackrel{\text{def}}{=} \begin{bmatrix} -3/2 \\ -3/2 \\ -3/2 \\ -3/2 \\ -1/2 \\ -1/2 \\ -1/2 \\ -1/2 \end{bmatrix} \cdot \|\boldsymbol{\mu}\| \cdot C, \quad \mathbf{r} \stackrel{\text{def}}{=} R \cdot \begin{bmatrix} 2 \\ -2 \\ -2 \\ 2 \\ -1 \\ -1 \\ -1 \\ -1 \end{bmatrix}, \quad (11)$$

556 where $R > 0$ and $C > 0$ are arbitrary scaling parameters. Both C and R and input to the score
 557 function are set different for each of the models, as indicated in Table 4. In particular, we set $R = \frac{7}{\|\boldsymbol{\mu}\|}$
 558 for both GAT and CAT such that: i) all γ_{ij} are distinguishable as we decrease $\|\boldsymbol{\mu}\|$; and ii) we avoid
 559 numerical instabilities in the implementation when computing the exponential of $R \times \|\boldsymbol{\mu}\|$ in order
 560 to obtain γ_{ij} (see Corollary 8), as the exponential of small or large values leads to under/overflow
 561 issues. As for C , we set $C = 1$ for GAT and $C = (p - q)/(p + 2q)$ for CAT such that we counteract
 562 the fact that the distance between classes shrink as we increase q , see Lemma 3:

Table 4: Parameters for the synthetic experiments.

Model	C	R	\mathbf{h}_i
GCN	0	—	—
GAT	1	$\frac{7}{\ \boldsymbol{\mu}\ }$	\mathbf{X}_i
CAT	$\frac{p-q}{p+2q}$	$\frac{7}{\ \boldsymbol{\mu}\ }$	$\frac{1}{ N_i^* } \sum_{k \in N_i^*} \mathbf{X}_k$

563 Regarding the data model, we set (as described in §6.1) $n = 10000$, $p = 0.5$, $\sigma = 0.1$, and
 564 $d = n / (5 \log^2(n))$. We set the slope of the LeakyReLU activation to $\beta = 1/5$ for the GAT and

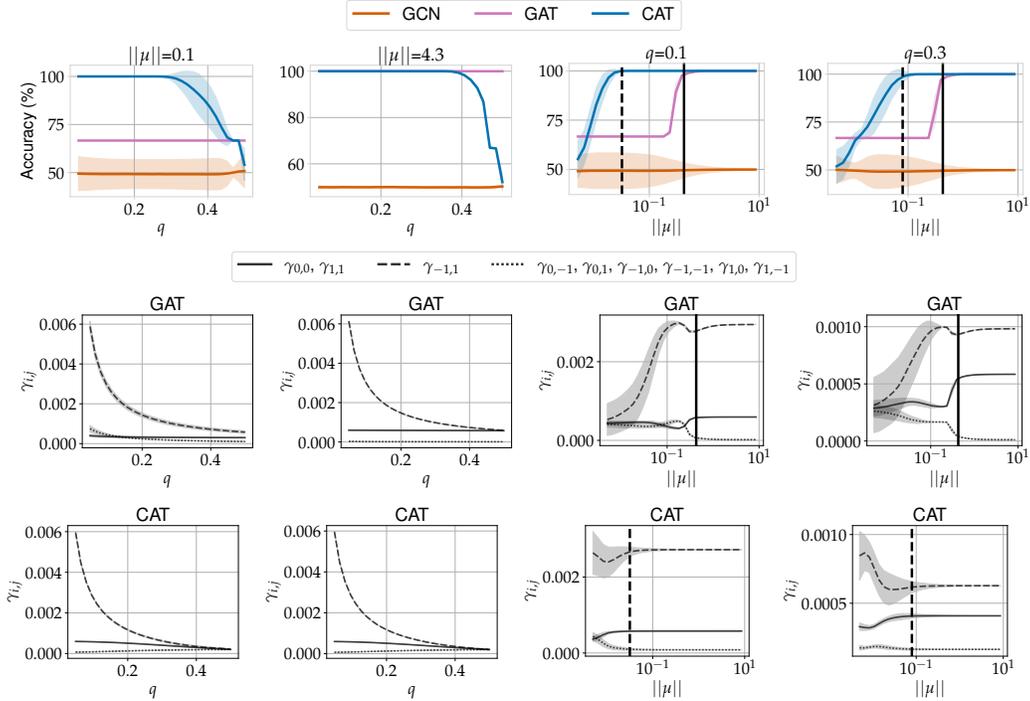


Figure 4: Synthetic data results. On the top row, we show the node classification, and in the following two rows we show the γ_{ij} values for GAT and CAT respectively. In the two left-most figures, we show how the results vary with the noise level q for $\|\mu\| = 0.1$ and $\|\mu\| = 4.3$. In the two right-most figures, we show how the results vary with the norm of the means $\|\mu\|$ for $q = 0.1$ and $q = 0.3$. We use two vertical lines to present the classification threshold stated in [Theorem 1](#) (solid line) and [Corollary 2](#) (dashed line).

565 $\beta = 0.01$ for CAT, such that the proof of [Corollary 8](#) is valid. As described in the main paper, to
566 assess the sensitivity to structural noise, we present the complete results for two sets of experiments.
567 First, we vary the noise level q between 0 and 0.5, fixing the mean vector μ . We test two values
568 of $\|\mu\|$: the first corresponds to the *easy* regime ($\|\mu\| = 10\sigma\sqrt{2\log n} \approx 4.3$) where classes are
569 far apart; and the second correspond to the *hard* regime ($\|\mu\| = \sigma = 0.1$) where the distance
570 between the clusters is small. In the second experiment we modify instead the distance between
571 the means, sweeping $\|\mu\|$ in the range $[\sigma/20, 20\sigma\sqrt{2\log n}]$ which corresponds to $[0.005, 8.58]$, and
572 thus covering the transition from the hard setting (small $\|\mu\|$) to the easy one (large $\|\mu\|$). In these
573 experiments, we fix q to 0.1 (low noise) and 0.3 (high noise).

574 **Results** are summarized in [Fig. 4](#). The top row contains the node classification performance for each
575 of the models (i.e., [Fig. 1](#)), the next two rows contain the γ_{ij} values for GAT and CAT respectively.
576 The two left-most columns of [Fig. 4](#) show the results for the hard and easy regimes, respectively, as
577 we vary the noise level q . In the hard regime, we observe that GAT is unable to achieve separation for
578 any value of q , whereas CAT achieves perfect classification when q is small enough. The gamma
579 plots help shed some light on this question. For GAT, we observe that the gammas represented with
580 the dotted and solid lines collapse for any value of q (see middle plot), while this does not happen for
581 CAT when the noise level is low (see bottom plot). This exemplifies the advantage of CAT over GAT
582 as stated in [Corollary 2](#). When the distance between the means is large enough, we see that GAT
583 achieves perfect results independently of q , as stated in [Theorem 1](#). We also observe that, in this case,
584 the gammas represented with the dotted and solid lines do not collapse for any value of q . In contrast,
585 as we increase q , CAT fails to satisfy the condition in [Corollary 2](#), and therefore achieves inferior
586 performance. We note that the low performance is due to the fact that all gammas collapse to the
587 same value for large noise levels.

588 For the second set of experiments (two right-most columns of [Fig. 1](#)), where we fix q and sweep
589 $\|\mu\|$, we observe that, for both values of q , there exists a transition in the accuracy of both GAT

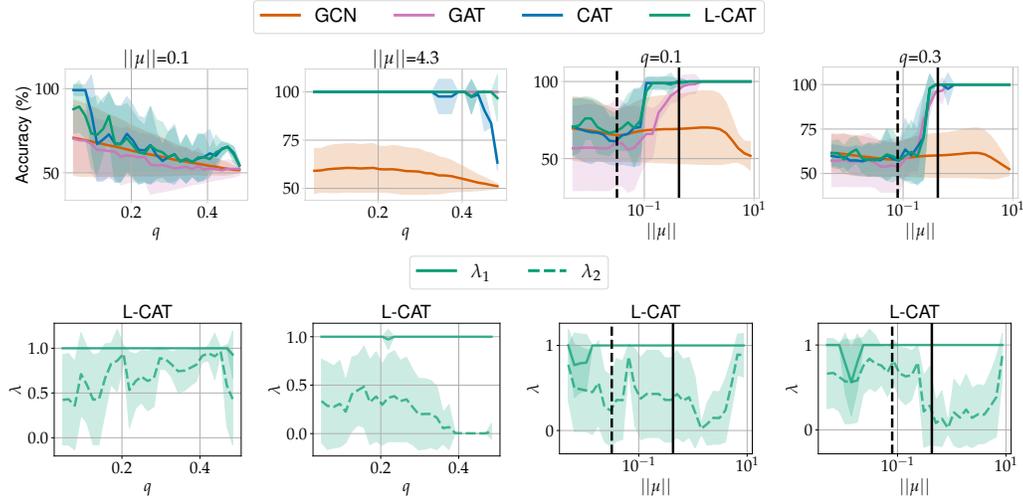


Figure 5: Synthetic data results learning C , λ_1 and λ_2 . On the top row, we show the node classification accuracy, and in the bottom row we show the learned values of λ_1 and λ_2 for L-CAT. In the two left-most figures, we show how the results vary with the noise level q for $\|\mu\| = 0.1$ and $\|\mu\| = 4.3$. In the two right-most figures, we show how the results vary with the norm of the means $\|\mu\|$ for $q = 0.1$ and $q = 0.3$. We use two vertical lines to present the classification threshold stated in [Theorem 1](#) (solid line) and [Corollary 2](#) (dashed line).

590 and CAT as a function of $\|\mu\|$. As shown in the main manuscript, GAT achieves perfect accuracy
 591 when the distance between the means satisfies the condition in [Theorem 1](#) (solid vertical line in
 592 [Fig. 1](#)). Moreover, we can see the improvement CAT obtains over GAT. Indeed, when $\|\mu\|$ satisfies
 593 the conditions of [Corollary 2](#) (dashed vertical line in [Fig. 1](#)), the classification threshold is improved.
 594 As we increase q , we see that the gap between the two vertical lines decreases, which means that the
 595 improvement decreases as q increments, exactly as stated in [Corollary 2](#). This transition from the
 596 hard regime to the easy regime is also observed in the gamma plots: we observe the largest difference
 597 in value between the different groups of lambdas for values of $\|\mu\|$ that satisfy the condition in
 598 [Theorem 1](#) (that is to the right of the vertical lines).

599 B.1 Other experiments

600 In the following, we extend the results for the synthetic data presented above. In particular, we aim to
 601 evaluate if L-CAT is able to achieve top performance regardless of the scenario. That is, we want to
 602 evaluate if L-CAT consistently performs at least as good as the best-performing model. We change
 603 the fixed-parameter setting of the previous section and, instead, we evaluate the performance of GCN,
 604 GAT, CAT and L-CAT when we learn the model-dependent parameters.

605 **Experimental setup.** We assume the same parametrization for the 1-layer GCN, GAT and CAT
 606 described in [Eq. 8](#) and [Eq. 10](#). For L-CAT, we add the parameters λ_1 and λ_2 , as indicated in [Eq. 7](#).
 607 We fix the parameters shared among the models, that is, \tilde{w} , S , b , r , and R , with the values indicated
 608 in [Eq. 11](#). Different from previous experiments, we now learn C and, for L-CAT, we also learn λ_1
 609 and λ_2 . We choose to fix part of the parameters (instead of learning them all) to keep the problem as
 610 similar as possible to the theoretical analysis we provided in [§4](#) and [Appendix A](#). If we instead learn
 611 all the parameters, it takes a single dimension of the features to be (close to) linearly separable to find
 612 a solution that achieves a similar performance regardless of the model, which hinders the analysis.
 613 This is a consequence of the probabilistic nature of the features. One way of solving this issue would
 614 be to make n big enough. Instead, we opt to have a fixed n and reduce the degrees of freedom of the
 615 models by fixing the parameters shared across all models. The rest of the experimental setup matches
 616 the one from [Appendix B](#). Additionally, we use the Adam optimizer [\[24\]](#) with a learning rate of 0.05,
 617 and we train for 100.

618 **Results** are summarized in [Fig. 5](#). The top row contains the node classification performance for
 619 every model, while the bottom row contains the learned values of λ_1 (solid line) and λ_2 (dashed line)

620 with L-CAT. The two left-most columns of Fig. 5 show the results for the hard and easy regimes,
 621 respectively, as we vary the noise level q . In the hard regime, we see rather noisy results. Still,
 622 the behaviour is similar to that of Fig. 4: the performance of CAT degrades as we increase q . We
 623 also observe that, on average, CAT outperforms GAT. In this case, we observe that L-CAT achieves
 624 similar performance as CAT, which can be explained by inspecting the learned values of lambda in
 625 the bottom row. We observe that $\lambda_1 = 1$ and $\lambda_2 \geq 0.5$ on average for all values of q . This indicates
 626 that L-CAT is closer to CAT than to GAT. When the distance between the means is large enough (i.e.,
 627 $\|\mu\| = 4.3$), we see that GAT achieves perfect results independently of q while the performance of
 628 CAT deteriorates with large values of q , the same trend as in Fig. 4. Remarkably, we observe that
 629 L-CAT also achieves perfect results independently of q . If we inspect the lambda values, we first see
 630 that $\lambda = 1$ for all q , thus the interpolation happens between CAT and GAT. Looking at the values of
 631 λ_2 , we observe that, for small values of q , λ_2 is pretty noisy, which is expected since any solution
 632 achieves perfect performance. Interestingly, we have that $\lambda_2 = 0$ for large values of q , with negligible
 633 variance. This indicates that L-CAT learns that it must behave like GAT in order to perform well.

634 For the second set of experiments (two right-most columns of Fig. 5), we fix q and sweep $\|\mu\|$ like
 635 we did in Fig. 4. Here, we observe a similar trend: for both values of q , there exists a transition in
 636 the accuracy of both GAT and CAT as a function of $\|\mu\|$. Yet once again, we observe that L-CAT
 637 consistently achieves a similar performance to the best-performing model in every situation.

638 C Dataset description

639 We present further details about the datasets used in our experiments, summarized in Table 5. All
 640 datasets are undirected (or transformed to undirected otherwise) and transductive.

641 The upper rows of the table refer to datasets used in §6.2 taken from the PyTorch Geometric
 642 framework.⁶ The following paragraphs present a short description of such datasets.

643 **Amazon Computers & Photos** are datasets taken from [34], in which nodes represent products, and
 644 edges indicate that the products are usually bought together. The node features are a Bag of Words
 645 (BoW) representation of the product reviews. The task is to predict the category of the products.

646 **GitHub** is a dataset introduced in [31], in which nodes correspond to developers, and edges indicate
 647 mutual follow relationship. Node features are embeddings extracted from the developer’s starred
 648 repositories and profile information (e.g., location or employer). The task is to infer whether a node
 649 relates to web or machine learning development.

650 **FacebookPagePage** is a dataset introduced in [31], where nodes are Facebook pages, and edges
 651 imply mutual likes between the pages. Nodes features are text embeddings extracted from the pages’
 652 description. The task consist on identifying the page’s category.

653 **TwitchEN** is a dataset introduced in [31]. Here, nodes correspond to Twitch gamers, and links reveal
 654 mutual friendship. Node features are an embedding of games liked, location, and streaming habits.
 655 The task is to infer if a gamer uses explicit content.

656 **Coauthor Physics & CS** are datasets introduced in [34]. In this case, nodes represent authors which
 657 are connected with an edge if they have co-authored a paper. Node features are BoW representations
 658 of the keywords of the author’s papers. The task consist on mapping each author to their corresponding
 659 field of study.

660 **DBLP** is a dataset introduced in [7] that represents a citation network. In this dataset, nodes represent
 661 papers and edges correspond to citations. Node features are BoW representations of the keywords of
 662 the papers. The task is to predict the research area of the papers.

663 **PubMed, Cora & CiteSeer** are citation networks introduced in [46]. Nodes represent documents,
 664 and edges refer to citations between the documents. Node features are BoW representations of the
 665 documents. The task is to infer the topic of the documents.

666 The bottom rows of Table 5 refer to the datasets from Open Graph Benchmark (OGB) [20]⁷ used in
 667 §6.3. We include a short description of them in the paragraphs below.

⁶<https://pytorch-geometric.readthedocs.io/en/latest/modules/datasets.html>

⁷<https://ogb.stanford.edu/docs/nodeprop>

668 **ogbn-arxiv** is a citation network of computer science papers in arXiv [40]. Nodes represent papers,
 669 and directed edges refer to citations among them. Node features are embeddings of the title and
 670 abstract of the papers. The task is to predict the research area of the nodes.

671 **ogbn-products** contains a co-purchasing network [6]. Nodes represent products, and links are present
 672 whenever two products are bought together. Node features are embeddings of a BoW representation
 673 of the product description. The task is to infer the category of the products.

674 **ogbn-mag** is a heterogeneous network formed from a subgraph of the Microsoft Academic Graph
 675 (MAG) [40]. Nodes can belong to one of these four types: authors, papers, institutions and fields of
 676 study. Moreover, directed edges belong to one of the following categories: “author is affiliated with
 677 an institution,” “author has written a paper,” “paper cites a paper,” and “paper belongs to a research
 678 area.” Only nodes that are papers contain node features, which are a text embedding of the document
 679 content. The task is to predict the venue of the nodes that are papers.

680 **ogbn-proteins** is a network whose nodes represent proteins and edges indicate different types of
 681 associations among them. This dataset does not contain node features. The tasks are to predict
 682 multiple protein functions, each of them being a binary classification problem.

Table 5: Dataset statistics. On the top part of the table, we show the datasets used in §6.2. On the
 bottom part of the table, we show the datasets used in §6.3.

Name	#Nodes	#Edges	Avg. degree	#Node feats.	#Edge feats.	#Tasks	Task Type
AmazonComp.	13,752	491,722	35.76	767	-	1	10-class clf.
AmazonPhoto	7,650	238,162	31.13	745	-	1	8-class clf.
GitHub	37,700	578,006	15.33	128	-	1	Binary clf.
FacebookP.	22,470	342,004	15.22	128	-	1	4-class clf.
CoauthorPh.	34,493	495,924	14.38	8415	-	1	5-class clf.
TwitchEN	7,126	77,774	10.91	128	-	1	Binary clf.
CoauthorCS	18,333	163,788	8.93	6805	-	1	15-class clf.
DBLP	17,716	105,734	5.97	1639	-	1	4-class clf.
PubMed	19,717	88,648	4.50	500	-	1	3-class clf.
Cora	2,708	10,556	3.90	1433	-	1	7-class clf.
CiteSeer	3,327	9,104	2.74	3703	-	1	6-class clf.
ogbn-arxiv	169,343	1,166,243	6.89	128	-	1	40-class clf.
ogbn-products	2,449,029	123,718,280	50.52	100	-	1	47-class clf.
ogbn-mag	1,939,743	21,111,007	18.61	128	4	1	349-class clf.
ogbn-proteins	132,534	79,122,504	597.00	-	8	112	Multi-task

683 D Real data experiments

684 D.1 Experimental details

685 **Computational resources.** We used CPU cores to run this set of experiments. In particular, for each
 686 trial, we used 2 CPU cores and up to 16 GB of memory. We ran the experiments in parallel using a
 687 shared cluster with 10000 CPU cores approximately.

688 **General experimental setup.** As mentioned in §6.2, we repeat all experiments 10 times, which
 689 correspond to 10 different random initialization of the parameters of the GNNs. In all cases, we
 690 choose the model parameters with the best validation performance during training. In order to run the
 691 experiments and collect the results, we used the GraphGym framework [47], which includes the data
 692 processing and loading of the datasets, as well as the evaluation and collection of the results. We split
 693 the datasets in 70 % training, 15 % validation, and 15 % test.

694 We cross-validate the number of message-passing layers in the network (2, 3, 4), as well as the
 695 learning rate ([0.01, 0.005]). Then, we report the results of the best validation error among the
 696 4 possible combinations. However, in practice we found the best performance always to use 4
 697 message-passing layers, and thus the only difference in configuration lies in the learning rate.

We use residual connections between the GNN layers, 4 heads in the attention models, and the Parametric ReLU (PReLU) [19] as the nonlinear activation function. We do not use batch normalization [22], nor dropout [35]. We use the Adam optimizer [24] with $\beta = (0.9, 0.999)$, and an exponential learning-rate scheduler with $\gamma = 0.998$. We train all the models for 2500 epochs. Importantly, we do not use weight decay, since this will bias the solution towards $\lambda_1 = 0$ and $\lambda_2 = 1$.

We use the Pytorch Geometric [13] implementation of L-CAT for all experiments, switching between models by properly by setting λ_1 and λ_2 . We parametrize λ_1 and λ_2 as free-parameters in log-space that pass through a sigmoid function—i.e., $\text{sigmoid}(10^x)$ —such that they are constrained to the unit interval, and they are learned quickly.

D.2 Additional results

Table 6 shows the results presented in the main paper (with the addition of a dense feed-forward network), while Table 7 presents the results for the remaining datasets, with smaller average degree.

If we focus on Table 7, we observe that all models perform equally well, yet in a few cases CAT and L-CAT are significantly better than the baselines—e.g., L-CATv2 in *CoauthorCS*, or L-CAT in *Cora*. Following a similar discussion as the one presented in the main paper, these results indicates that L-CAT achieves similar or better performance than baseline models and thus, should be the preferred architecture.

Competitive performance without the graph. We also include in Tables 6 and 7 the performance of a feed-forward network, referred to as Dense (first row). Note that the only data available to this model are the node features, and thus no graph information is provided. Therefore, we should expect a significant drop in performance, which indeed happens for some datasets such as *Amazon Computers* ($\approx 7\%$ drop), *FacebookPagePage* ($\approx 20\%$ drop), *DBLP* ($\approx 9\%$ drop) and *Cora* ($\approx 14\%$ drop). Still, we found that for other commonly used datasets the performance is similar, e.g., *Coauthor Physics* and *PubMed*; or it is even better *CoauthorCS*. These results manifest the importance of a proper benchmarking, and of carefully considering the datasets used to evaluate GNN models.

Table 6: Test accuracy (%) of the considered convolution and attention models for different datasets (sorted by their average node degree), and averaged over ten runs. Bold numbers are statistically different to their baseline model ($\alpha = 0.05$). Best average performance is underlined.

Dataset	<i>Amazon Computers</i>	<i>Amazon Photo</i>	<i>GitHub</i>	<i>Facebook PagePage</i>	<i>Coauthor Physics</i>	<i>TwitchEN</i>
Avg. Deg.	35.76	31.13	15.33	15.22	14.38	10.91
Dense	83.73 \pm 0.34	91.74 \pm 0.46	81.21 \pm 0.30	75.89 \pm 0.66	95.41 \pm 0.14	56.26 \pm 1.74
GCN	<u>90.59 \pm 0.36</u>	<u>95.13 \pm 0.57</u>	84.13 \pm 0.44	94.76 \pm 0.19	96.36 \pm 0.10	57.83 \pm 1.13
GAT	89.59 \pm 0.61	94.02 \pm 0.66	83.31 \pm 0.18	94.16 \pm 0.48	96.36 \pm 0.10	57.59 \pm 1.20
CAT	90.58 \pm 0.40	94.77 \pm 0.47	84.11 \pm 0.66	94.71 \pm 0.30	96.40 \pm 0.10	58.09 \pm 1.61
L-CAT	90.34 \pm 0.47	94.93 \pm 0.37	84.05 \pm 0.70	94.81 \pm 0.25	96.35 \pm 0.10	57.88 \pm 2.07
GATv2	89.49 \pm 0.53	93.47 \pm 0.62	82.92 \pm 0.45	93.44 \pm 0.30	96.24 \pm 0.19	57.70 \pm 1.17
CATv2	90.44 \pm 0.46	94.81 \pm 0.55	84.10 \pm 0.88	94.27 \pm 0.31	96.34 \pm 0.12	57.99 \pm 2.02
L-CATv2	90.33 \pm 0.44	94.79 \pm 0.61	<u>84.31 \pm 0.59</u>	94.44 \pm 0.39	96.29 \pm 0.13	57.89 \pm 1.53

E Open Graph Benchmark experiments

E.1 Experimental details

Computational resources. For this set of experiments, we had at our disposal a set of 16 Tesla V100-SXM GPUs with 160 CPU cores, shared among the rest of the department.

Statistical significance. For each CAT and L-CAT model, we highlight significant improvements according to a two-sided paired t-test ($\alpha = 0.05$), with respect to its corresponding baseline model. For example, for L-CATv2 with 8 heads we perform the test with respect to GATv2 with 8 heads.

General experimental setup. As mentioned in §6.3, we repeat all experiments with OGB datasets 5 times. In all cases, we choose the model parameters with the best validation performance during

Table 7: Test accuracy (%) of the considered convolution and attention models for different datasets (sorted by their average node degree), and averaged over ten runs. Bold numbers are statistically different to their baseline model ($\alpha = 0.05$). Best average performance is underlined.

Model	<i>CoauthorCS</i>	<i>DBLP</i>	<i>PubMed</i>	<i>Cora</i>	<i>CiteSeer</i>
Avg. Deg.	8.93	5.97	4.5	3.9	2.74
Dense	<u>94.88 ± 0.21</u>	75.46 ± 0.27	88.13 ± 0.33	72.75 ± 1.72	73.02 ± 1.01
GCN	93.85 ± 0.23	84.18 ± 0.40	88.50 ± 0.18	<u>86.68 ± 0.78</u>	<u>75.76 ± 1.09</u>
GAT	93.80 ± 0.38	84.15 ± 0.39	<u>88.62 ± 0.18</u>	85.95 ± 0.95	75.40 ± 1.43
CAT	93.70 ± 0.31	84.10 ± 0.29	<u>88.58 ± 0.25</u>	85.85 ± 0.79	75.64 ± 0.91
L-CAT	93.65 ± 0.23	84.13 ± 0.26	88.45 ± 0.32	86.66 ± 0.87	75.04 ± 1.12
GATv2	93.19 ± 0.64	<u>84.33 ± 0.18</u>	88.52 ± 0.27	85.65 ± 1.01	75.14 ± 1.20
CATv2	93.51 ± 0.34	84.15 ± 0.41	88.54 ± 0.29	85.50 ± 0.94	74.68 ± 1.30
L-CATv2	93.65 ± 0.20	84.31 ± 0.31	88.48 ± 0.24	85.75 ± 0.72	75.04 ± 1.30

732 training. Moreover, when we show the results without specifying the number of heads, we take the
 733 model with the best validation error among the two models with 1 and 8 heads.

734 We use the same implementation of L-CAT for all experiments, switching between models by properly
 735 setting λ_1 and λ_2 . Experiments on *arxiv*, *mag*, *products* use a version of L-CAT implemented in
 736 Pytorch Geometric [13]. Experiments on *proteins* use a version of L-CAT implemented in DGL [41].
 737 We parametrize λ_1 and λ_2 as free-parameters in log-space that pass through a sigmoid function—i.e.,
 738 $\text{sigmoid}(10^x)$ —such that they are constrained to the unit interval, and they are learned quickly.

739 **ArXiv.** As described in §6.3, we use the example code from the OGB framework [20]. The network
 740 is composed of 3 GNN layers with a hidden size of 128. We use batch normalization [22] and a
 741 dropout [35] of 0.5 between the GNN layers, and Adam [24] with a learning rate of 0.01. We use the
 742 ReLU as activation function. For the initial experiments, we train for 1500 epochs, while we train for
 743 500 epochs for the noise experiments in §6.3.1. This is justified given the convergence plots in Fig. 2.

744 **MAG.** We adapted the official code from [8]. The network is composed of 2 layers with 128 hidden
 745 channels. This time, we use layer normalization [2] and a dropout of 0.5 between the layers. Again,
 746 we use ReLU as the activation function, and add residual connections to the network. As with *arxiv*,
 747 we use Adam [24] with learning rate 0.01. We set a batch size of 20000 and train for 100 epochs.

748 **Products.** We use the same setup as [8], with a network of 3 GNN layers and 128 hidden dimensions.
 749 We apply residual connections once again, with a dropout [35] of 0.5 between layers. This time, we
 750 use ELU as the activation function. The batch size is set to 256. Adam [24] is again the optimizer in
 751 use, this time with a learning rate of 0.001. We train for 100 epochs, although we apply early stopping
 752 whenever the validation accuracy stops increasing for more than 10 epochs. Note the training split of
 753 this dataset only contains 8% of the data.

754 **Proteins.** We follow once more the setup of [8]. The network we use has 6 GNN layers of hidden
 755 size 64. Dropout [35] is set to 0.25 between layers, with an input dropout of 0.1. At the beginning
 756 of the network, we place a linear layer followed by a ReLU activation to encode the nodes, and a
 757 linear layer at the end of the network to predict the class. Moreover, we use batch normalization [22]
 758 between layers and ReLU as the activation function. We train the model for 1200 epochs at most,
 759 with early stopping after not improving for 10 epochs.

760 E.2 Additional results

761 We show in Tables 8 to 10 the results of the main paper for the *arxiv*, *mag*, *products* datasets,
 762 respectively, without selecting the best configuration for each type of model. That is, we show the
 763 results for both number of heads. Note that we already show the full table of results for the *protein*
 764 datasets in the main paper (Table 3). All the trends discussed in the main paper hold.

765 **Extrapolation ablation study.** Due to page constraints, these results were not added to the main
 766 paper. Here, we study two questions. First, how important are λ_1 and λ_2 in the formulation of L-CAT
 767 (Eq. 7)? For the sake of completeness, the second question we attempt to answer here is whether we

Table 8: Test accuracy on the *arxiv* dataset for attention models using 1 head and 8 heads.

	GCN	GAT	CAT	L-CAT	GATv2	CATv2	L-CATv2
1h	71.58 ± 0.19	71.58 ± 0.15	72.04 ± 0.20	72.00 ± 0.11	71.70 ± 0.14	72.02 ± 0.08	71.96 ± 0.21
8h	–	71.63 ± 0.11	72.14 ± 0.20	71.98 ± 0.08	71.72 ± 0.24	71.76 ± 0.14	71.91 ± 0.16

Table 9: Test accuracy on the *mag* dataset for attention models using 1 head and 8 heads.

	GCN	GAT	CAT	L-CAT	GATv2	CATv2	L-CATv2
1h	<u>32.77 ± 0.36</u>	32.35 ± 0.24	31.98 ± 0.46	32.47 ± 0.38	32.76 ± 0.18	32.43 ± 0.22	32.68 ± 0.50
8h	–	32.15 ± 0.31	31.58 ± 0.22	32.49 ± 0.21	<u>32.85 ± 0.21</u>	32.34 ± 0.18	32.38 ± 0.28

Table 10: Test accuracy on the *products* dataset for attention models using 1 head and 8 heads.

	GCN	GAT	CAT	L-CAT	GATv2	CATv2	L-CATv2
1h	74.12 ± 1.20	<u>78.53 ± 0.91</u>	77.38 ± 0.36	77.19 ± 1.11	73.81 ± 0.39	74.81 ± 1.12	76.37 ± 0.92
8h	–	<u>78.23 ± 0.25</u>	76.63 ± 1.15	76.56 ± 0.45	76.40 ± 0.71	75.20 ± 0.92	74.70 ± 0.28

768 can obtain similar performance by just interpolating between GCN and GAT (fixing $\lambda_2 = 0$)? Note
 769 that we theoretically showed in §§4 and 6.1 that CAT fills up a gap between GCN and GAT, making
 770 it preferable in certain settings.

771 To this end, we repeat the experiments for network-initialization robustness in §6.3.2, since they
 772 showed to be the best ones to tell apart the performance across models. We include three additional
 773 models: GCN-GAT, which interpolates between GCN and GAT (or GATv2) by learning λ_1 and fixing
 774 $\lambda_2 = 0$; CAT- λ_1 which interpolates between GCN and CAT by learning λ_1 and fixing $\lambda_2 = 1$; and
 775 CAT- λ_2 , which interpolates between GAT and CAT by learning λ_2 and fixing $\lambda_1 = 1$.

776 Results using GAT and shown in Table 11, and using GATv2 in Table 12. We can observe that
 777 GCN-GAT obtains results in between GCN and GAT for all settings, despite being able to interpolate
 778 between both layers in each of the six layers of the network. Regarding learning λ_1 and λ_2 , we can
 779 observe that there is a clear difference between learning boths (L-CAT), and learning a single one.
 780 For both attention models, CAT- λ_1 obtains better results than CAT- λ_2 in all settings, but *uniform*
 781 with 8 heads. Still, the results of both variants are substantially worse than those of L-CAT in all
 782 cases, *demonstrating the importance of learning to interpolate between the three layer types.*

Table 11: Test accuracy on the *proteins* dataset for GCN [25] and GAT [38] attention models using two network initializations, and two numbers of heads (1 and 8).

	GCN	GCN-GAT	GAT	CAT	L-CAT	CAT- λ_1	CAT- λ_2
<i>uniform</i> initialization							
1h	61.08 ± 2.86	70.44 ± 1.56	59.73 ± 4.04	74.19 ± 0.72	<u>77.77 ± 1.44</u>	71.97 ± 3.78	73.55 ± 1.36
8h	–	68.51 ± 0.91	72.23 ± 3.20	73.60 ± 1.27	<u>78.85 ± 1.76</u>	76.43 ± 2.47	72.76 ± 2.79
<i>normal</i> initialization							
1h	<u>80.10 ± 0.61</u>	66.51 ± 3.23	66.38 ± 7.76	73.26 ± 1.84	78.06 ± 1.40	76.77 ± 1.91	73.39 ± 1.25
8h	–	69.93 ± 1.93	79.08 ± 1.64	74.67 ± 1.29	<u>79.63 ± 0.79</u>	78.86 ± 1.07	73.32 ± 1.15

Table 12: Test accuracy on the *proteins* dataset for GCN [25] and GATv2 [8] attention models using two network initializations, and two numbers of heads (1 and 8).

	GCN	GCN-GATv2	GATv2	CATv2	L-CATv2	CATv2- λ_1	CATv2- λ_2
<i>uniform</i> initialization							
1h	61.08 ± 2.86	69.69 ± 1.59	59.85 ± 3.05	64.32 ± 2.61	79.08 ± 1.06	63.24 ± 1.55	73.41 ± 0.34
8h	–	69.94 ± 1.62	75.21 ± 1.80	74.16 ± 1.45	<u>78.77 ± 1.09</u>	77.61 ± 1.32	73.96 ± 1.27
<i>normal</i> initialization							
1h	<u>80.10 ± 0.61</u>	68.54 ± 1.63	69.13 ± 9.48	74.33 ± 1.06	79.07 ± 1.09	78.41 ± 0.93	74.07 ± 1.17
8h	–	68.71 ± 1.96	78.65 ± 1.61	73.40 ± 0.62	<u>79.30 ± 0.55</u>	78.76 ± 1.41	73.22 ± 0.77

783 References

- 784 [1] T.W. Anderson. *An introduction to multivariate statistical analysis*. John Wiley & Sons, 2003.
- 785 [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint*
786 *arXiv:1607.06450*, 2016.
- 787 [3] Aseem Baranwal, Kimon Fountoulakis, and Aukosh Jagannath. Graph convolution for semi-
788 supervised classification: Improved linear separability and out-of-distribution generalization. In
789 *International Conference on Machine Learning (ICML)*. PMLR, 2021.
- 790 [4] Aseem Baranwal, Kimon Fountoulakis, and Aukosh Jagannath. Effects of graph convolutions
791 in deep networks. *arXiv preprint arXiv:2204.09297*, 2022.
- 792 [5] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius
793 Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan
794 Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint*
795 *arXiv:1806.01261*, 2018.
- 796 [6] K. Bhatia, K. Dahiya, H. Jain, P. Kar, A. Mittal, Y. Prabhu, and M. Varma. The extreme
797 classification repository: Multi-label datasets and code, 2016. URL [http://manikvarma.org/
798 downloads/XC/XMLRepository.html](http://manikvarma.org/downloads/XC/XMLRepository.html).
- 799 [7] Aleksandar Bojchevski and Stephan Günnemann. Deep gaussian embedding of graphs: Unsu-
800 pervised inductive learning via ranking. *arXiv preprint arXiv:1707.03815*, 2017.
- 801 [8] Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? In
802 *International Conference on Learning Representations (ICLR)*, 2022.
- 803 [9] Dan Busbridge, Dane Sherburn, Pietro Cavallo, and Nils Y Hammerla. Relational graph
804 attention networks. *arXiv preprint arXiv:1904.05811*, 2019.
- 805 [10] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal
806 neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*
807 (*NeurIPS*), 33, 2020.
- 808 [11] Yash Deshpande, Subhabrata Sen, Andrea Montanari, and Elchanan Mossel. Contextual
809 stochastic block models. *Advances in Neural Information Processing Systems (NeurIPS)*, 31,
810 2018.
- 811 [12] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs.
812 *arXiv preprint arXiv:2012.09699*, 2020.
- 813 [13] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric.
814 In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- 815 [14] Kimon Fountoulakis, Amit Levi, Shenghao Yang, Aseem Baranwal, and Aukosh Jagannath.
816 Graph attention retrospective. *arXiv preprint arXiv:2202.13060*, 2022.
- 817 [15] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl.
818 Neural message passing for quantum chemistry. In *International Conference on Machine*
819 *Learning (ICML)*, 2017.
- 820 [16] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward
821 neural networks. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*,
822 2010.
- 823 [17] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large
824 graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017.
- 825 [18] William L. Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods
826 and applications. *IEEE Data Eng. Bull.*, 40, 2017.
- 827 [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers:
828 Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE*
829 *international conference on computer vision*, 2015.
- 830 [20] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele
831 Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs.
832 *Advances in Neural Information Processing Systems (NeurIPS)*, 33:22118–22133, 2020.
- 833 [21] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. Heterogeneous graph transformer. In
834 *The Web Conference (WWW)*, 2020.

- 835 [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training
836 by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*.
837 PMLR, 2015.
- 838 [23] Dongkwan Kim and Alice Oh. How to find your friendly neighborhood: Graph attention design
839 with self-supervision. In *International Conference on Learning Representations (ICLR)*, 2021.
- 840 [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Inter-
841 national Conference on Machine Learning (ICML)*, 2015. URL [http://arxiv.org/abs/
842 1412.6980](http://arxiv.org/abs/1412.6980).
- 843 [25] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional
844 networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- 845 [26] Boris Knyazev, Graham W Taylor, and Mohamed Amer. Understanding attention and gen-
846 eralization in graph neural networks. *Advances in Neural Information Processing Systems
847 (NeurIPS)*, 32, 2019.
- 848 [27] Devin Kreuzer, Dominique Beaini, Will Hamilton, Vincent Létourneau, and Prudencio Tossou.
849 Rethinking graph transformers with spectral attention. *Advances in Neural Information Pro-
850 cessing Systems (NeurIPS)*, 34, 2021.
- 851 [28] John Boaz Lee, Ryan A Rossi, Sungchul Kim, Nesreen K Ahmed, and Eunye Koh. Attention
852 models in graphs: A survey. *ACM Transactions on Knowledge Discovery from Data (TKDD)*,
853 13, 2019.
- 854 [29] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen,
855 Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural
856 networks. In *AAAI conference on artificial intelligence*, volume 33, 2019.
- 857 [30] P. Rigollet and J.-C. Hütter. High dimensional statistics. *Lecture notes for course 18S997*, 813:
858 814, 2015.
- 859 [31] Benedek Rozemberczki, Carl Allen, and Rik Sarkar. Multi-scale attributed node embedding.
860 *Journal of Complex Networks*, 2021.
- 861 [32] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini.
862 The graph neural network model. *IEEE transactions on neural networks*, 20, 2008.
- 863 [33] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-
864 Rad. Collective classification in network data. *AI magazine*, 29(3), 2008.
- 865 [34] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann.
866 Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.
- 867 [35] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov.
868 Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine
869 Learning Research*, 15(56):1929–1958, 2014.
- 870 [36] Kiran K Thekumparampil, Chong Wang, Sewoong Oh, and Li-Jia Li. Attention-based graph
871 neural network for semi-supervised learning. *arXiv preprint arXiv:1803.03735*, 2018.
- 872 [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
873 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information
874 Processing Systems (NeurIPS)*, 30, 2017.
- 875 [38] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua
876 Bengio. Graph attention networks. In *International Conference on Learning Representations
877 (ICLR)*, 2018.
- 878 [39] Guangtao Wang, Rex Ying, Jing Huang, and Jure Leskovec. Multi-hop attention graph neural
879 networks. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- 880 [40] Kuansan Wang, Zhihong Shen, Chiyuan Huang, Chieh-Han Wu, Yuxiao Dong, and Anshul
881 Kanakia. Microsoft academic graph: When experts are not enough. *Quantitative Science
882 Studies*, 2020.
- 883 [41] Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma,
884 Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang.
885 Deep graph library: A graph-centric, highly-performant package for graph neural networks.
886 *arXiv preprint arXiv:1909.01315*, 2019.

- 887 [42] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heteroge-
888 neous graph attention network. In *The World Wide Web Conference (WWW)*, 2019.
- 889 [43] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. Graph neural networks in recom-
890 mender systems: a survey. *ACM Computing Surveys (CSUR)*, 2020.
- 891 [44] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A
892 comprehensive survey on graph neural networks. *IEEE transactions on neural networks and*
893 *learning systems*, 32, 2020.
- 894 [45] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
895 networks? In *International Conference on Learning Representations (ICLR)*, 2019.
- 896 [46] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning
897 with graph embeddings. In *International conference on machine learning*, 2016.
- 898 [47] Jiaxuan You, Zhitao Ying, and Jure Leskovec. Design space for graph neural networks. *Advances*
899 *in Neural Information Processing Systems (NeurIPS)*, 33, 2020.
- 900 [48] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. Graph
901 transformer networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 32,
902 2019.