

Towards Semantic Understanding of GNN Layers Embedding with Functional-Semantic Activation Mapping

Kislay Raj^{a,*} and Alessandra Mileo^a

^a *INSIGHT Centre for Data Analytics & School of Computing, Dublin City University, Ireland, Dublin*
E-mail: kislay.raj2@mail.dcu.ie

Abstract.

Graph Neural Networks (GNNs) have demonstrated significant potential in learning representations from complex graph-structured data. However, current explainability approaches remain predominantly focused on identifying locally relevant sub-graphs (instance-level explanations), while global model interpretability (model-level explanations) is still an open research challenge. In addition, the impact of the GNN structure on the quality of the deep representation is poorly understood, with optimal layer configuration still being determined empirically through trial and error. In this paper, we extend our previous work on Functional-Semantic Activation Mapping (FSAM) to investigate how changing the number of GNN layers affects the quality and performance of the deep representation. Through experiments on multiple datasets, we observe that while adding layers may enhance accuracy, it does not consistently lead to improved semantic representations as, in some cases, performance may increase while semantic quality declines, suggesting correct predictions for incorrect reasons. FSAM layer-wise activation tracking allowed us to track neuron activations across layers, revealing that deeper layers can reduce neuron specialisation and lead to class misclassifications. Community analysis further indicates that certain misclassified classes share neurons in overlapping communities, highlighting a loss of class-specific representations at greater depths. Our findings demonstrate a critical trade-off that increased depth can compromise interpretability without commensurate gains in meaningful semantic learning. Instead of chasing accuracy alone, we need frameworks assessing whether models are learning coherent patterns; this is where FSAM proves invaluable, as it is emerging as a vital model-level diagnostic tool for architectural analysis.

Keywords: Explainable AI, Graph Neural Network, Graph Analysis, Neuro-Symbolic AI

1. Introduction

GNNs [7, 21, 30] have shown remarkable performance across node classification, link prediction, and graph classification tasks. GNNs leverage structural information and node features to capture complex relationships within a graph. However, explaining GNN predictions remains a challenge due to the complex topological nature of graphs and how this is represented in GNN embeddings. Unlike traditional neural networks, GNNs operate on graph structures, which might suggest better interpretability, but understanding how these relationships are learned within the layers remains ambiguous. In current research on GNN explainability, most local methods [31] focus on generating small subgraphs and identifying which nodes and edges contribute to a specific prediction. However, they do

*Corresponding author. E-mail: kislay.raj2@mail.dcu.ie.

1 not explicitly show how information is processed within the network layers. While useful, they fail to comprehen- 1
2 sively understand how the GNN behaves across different layers and do not offer a global understanding of model 2
3 behaviour. In our previous work [14], we introduced FSAM, which explains how the entire model behaves across 3
4 different layers. This makes it a more assertive, more interpretable and more structured approach than traditional 4
5 local explainability methods, which are limited to isolated explanations without deeper insights into network-wide 5
6 behaviour. In this extended work, we explore a key question in GNN design: How much does adding more layers 6
7 enhance the model’s ability to represent network behaviour? And does improved performance always mean better 7
8 representations? 8

9 The over-smoothing has been well-studied in the literature and is a well-known problem in GNNs [24, 8, 15], which 9
10 occurs when we add more layers of information to a GNN architecture. Our findings indicate that over-smoothing 10
11 reduces FSAM quality, as evidenced by the degradation of the model’s ability to represent meaningful, class-specific 11
12 features across layers. Instead of merely confirming that over-smoothing occurs, FSAM provides a better-structured 12
13 way to detect and quantify the effects of over-smoothing at the neuron level. It tracks when and where neurons begin 13
14 to lose their class-specific activations, offering more profound insight into the impact of model depth on representa- 14
15 tion quality. In this sense, FSAM is more likely to be a diagnostic tool for global-level model behaviour, indicating 15
16 where over-smoothing might occur. Furthermore, our research has practical implications. We have observed cases 16
17 where GNN performance improves without a corresponding enhancement in FSAM quality. This suggests that the 17
18 model may make correct predictions, but not necessarily for the right reasons, as it could rely on less meaningful, 18
19 over-smoothed features. Therefore, FSAM offers an interpretability-driven analysis of over-smoothing, providing 19
20 insights into the reasons behind the model’s predictions. Our findings elucidate the tradeoff between model depth and 20
21 interpretability, demonstrating empirically how excessive layering can degrade semantic coherence while maintain- 21
22 ing superficial accuracy metrics. The FSAM framework emerges as an essential diagnostic tool, offering researchers 22
23 the unprecedented capability to (i) quantify the progressive loss of neuron specialisation across layers, (ii) identify 23
24 where correct classifications stem from flawed reasoning patterns, and (iii) establish optimal depth thresholds before 24
25 semantic collapse occurs, establishing a new paradigm for assessing both what GNNs predict and how they derive 25
26 these predictions—a crucial distinction for deploying graph networks in high-stakes real world applications. 26

27 Since this paper relies on capturing the GNN’s behaviour through activation analysis with FSAM, our first contribu- 27
28 tion is to extend FSAM validation beyond our previous experiments on CORA [19] and CiteSeer [11]. To 28
29 achieve this, we conduct additional experiments on four datasets: PubMed [3], Amazon Computers [10], Amazon 29
30 Photos [10], and Coauthor [20]. These datasets, with their distinct topological complexities, allow us to compre- 30
31 hensively evaluate FSAM’s approach and determine how well the resulting activation graph reflects the GNN’s 31
32 behaviour and how effectively the network learns the semantic structure of the input data. 32

33 The contributions of this work can be summarized as follows. Firstly, we extend the FSAM approach by conduct- 33
34 ing experiments on a broader range of datasets to validate that the activation analysis and graphs generated by FSAM 34
35 consistently reflect the network behaviour. This includes community analysis in different datasets that demonstrates 35
36 the ability of FSAM to capture the semantic structure between classes reliably. Secondly, we extend our experimen- 36
37 tal analysis to confirm that the functional activation graph generated by FSAM aligns with the network’s behaviour 37
38 as the number of layers changes. By testing networks with different depths (from 1 to 4 layers) and comparing the 38
39 correlation between misclassifications and class similarity, we show that improvements in network accuracy are re- 39
40 flected in the FSAM graph, and the FSAM structure also captures any decline in accuracy. This analysis emphasizes 40
41 FSAM’s ability to represent network behaviour across different layer configurations accurately. Third, we conduct a 41
42 detailed layer-by-layer analysis to assess how different GNN layer configurations affect the model’s performance in 42
43 node classification tasks. Specifically, we examine how varying the number of layers influences FSAM and the cor- 43
44 responding community structure and verify whether improvements in accuracy align with better FSAM graphs and, 44
45 on the other hand, decreases in accuracy correlate with a decline in FSAM quality. This analysis demonstrates that 45
46 while additional layers may enhance performance, deeper layers can lead to over-smoothing and neuron activation 46
47 overlap, ultimately diminishing the model’s ability to differentiate between classes. As part of our comprehensive 47
48 analysis, we also identify a few interesting cases where accuracy improves without a corresponding improvement 48
49 in FSAM’s semantic quality. These instances reveal situations where the GNN achieves better predictions, but not 49
50 necessarily due to a better embedding the semantic structure in the input data. It highlights FSAM’s potential in 50
51 identifying cases where a model makes accurate predictions for the wrong reasons. 51

2. State of the Art

The interest in neurosymbolic AI has steadily increased, driven by the need for interpretable and accountable machine learning systems, especially in domains requiring transparent decision-making. Research has focused on integrating neural learning with symbolic reasoning, an essential step for enhancing the explainability of deep learning models. This integration is crucial for high-stakes domains where accuracy and interpretability are essential. GNNs have shown exceptional performance in handling graph-structured data across a range of fields, such as social networks [27], molecular structures [5], and citation networks [25]. However, despite their success in learning complex relationships, GNNs remain largely opaque regarding how specific predictions are made, particularly when compared to models in other domains like image and text analysis. The challenge lies in interpreting the internal representations learned by GNNs, particularly about prior knowledge.

Most existing methods for GNN explainability focus on *local explanations*, identifying key input features, nodes, or edges influencing individual predictions. These techniques are broadly divided into several categories: **Gradient/Feature-based methods** [13], which use gradient information or hidden feature map values to assess feature importance; **Perturbation-based methods** [28], which modify graph structures and monitor how these perturbations affect model outputs; **Decomposition methods** [17, 13], which break down the prediction score into contributions from different neurons or layers, propagating these contributions backwards through the network; and **Surrogate methods** [6, 22], which train interpretable models to approximate the GNN's behaviour by sampling the input graph's neighbourhood and constructing an explanation based on the simplified model.

Although these methods provide valuable insights, they predominantly focus on instance-level predictions, failing to capture the GNN's global behaviour or how information is processed through the network layers. Local methods often highlight specific features without a comprehensive view of the decision-making process, which is essential for understanding the model's behaviour about prior knowledge and domain-specific concepts.

It is worth stating that the use of *global explanations* is less researched in the context of GNN. One of the methods is **XGNN** [32], which creates synthetic graphs tailored to class predictions to explain the behaviour of the GNN. However, XGNN's assumption that a single synthetic graph can represent an entire class of graphs is unrealistic, as many relationships exist within real-world datasets. Empirical studies and follow-up methods have pointed out this shortcoming. For instance, in molecular graph classification, several different substructures (motifs) can cause a molecule to be mutagenic [1]. XGNN tends to identify one dominant motif and fails to identify other motifs, thus lacking multi-modal explanations. Note that generating one graph per class erases any combinatorial aspect that the GNN could have learned; if the model has a concept for a class as a conjunction or disjunction of several patterns, the XGNN explanation will be partial. Another concrete example comes from chemical ring structures: **MAGE** [29] found that prior model-level explainers like XGNN often fail to identify certain valid substructures (e.g. rings in molecules), leading to questionable interpretability. It happens because XGNN builds graphs edge-by-edge ("atom-by-atom"), which can struggle to capture a complex motif that requires adding a set of edges together (closing a ring). XGNN's inability to capture multiple modes or diverse substructures within the same class means its explanations can be incomplete or biased toward one pattern. In contrast, studies show that more nuanced approaches can uncover a richer set of class-specific patterns. In recent research, **GLGExplainer** [1] addresses these gaps by providing global explanations as Boolean combinations of learned graphical concepts. Unlike XGNN, **GLGExplainer** does not rely on synthetic prototypes but instead aggregates local explanations into interpretable concepts, which are then combined into logic formulas. as it learns concepts purely from local explanations without grounding them in human-interpretable semantics (e.g., chemical functional groups in molecules). However, its reliance on local explanations, discrete clustering, and lack of concept grounding limit its ability to align with the GNN's internal reasoning fully.

Although these approaches provide some information about the final predictions, they do not explain how the intermediate layers participate in the learned representations as they are not suitable for explaining the relationship between the internal structure of the model and prior knowledge or domain knowledge, which limits the ability of users to trust and understand the model's decisions. Furthermore, the related work in SOTA is presented in Table. 1.

Our previous work addresses this gap by introducing the **FSAM** approach, which provides a global explanation of GNNs by extracting deep representations in the form of semantic graphs. FSAM focuses on capturing the global structure of the GNN, along with the semantic relationships between neurons across different layers. This method

Table 1

Comparison of existing explainability methods for GNNs. The columns represent various aspects of each method: - **TYPE**: Indicates whether the method is **instance-level** (Local) or **model-level** (Global). - **LEARNING**: Specifies whether the method uses **backward** or **forward** propagation for explanation. - **TASK**: The tasks the method is designed to explain, such as **GC** (Graph Classification) or **NC** (Node Classification). - **TARGET**: The target to be explained, including **N** (node), **E** (edge), **NF** (node features), or **Subgraph**. - **Layer wise interpretability**: Describes whether the method tracking semantic changes across layers (\checkmark) or (\times). Abbreviations for task, target, layer wise interpretability are explained in the caption.

Method	TYPE	LEARNING	TASK	TARGET	LAYER-WISE INTERPRETABILITY
SA [2, 13]	Instance-level	\times	GC/NC	N/E/NF	\times
Guided BP [2]	Instance-level	\times	GC/NC	N/E/NF	\times
CAM [13]	Instance-level	\times	GC	N	\times
Grad-CAM [13]	Instance-level	\times	GC	N	\times
GNNExplainer [28]	Instance-level	\checkmark	GC/NC	E/NF	\times
PGExplainer [9]	Instance-level	\checkmark	GC/NC	E	\times
GraphMask [16]	Instance-level	\checkmark	GC/NC	E	\times
ZORRO [4]	Instance-level	\times	GC/NC	N/NF	\times
Causal Screening [23]	Instance-level	\times	GC/NC	E	\times
SubgraphX [31]	Instance-level	\checkmark	GC/NC	Subgraph	\times
LRP [2, 18]	Instance-level	\times	GC/NC	N	\times
Excitation BP [13]	Instance-level	\times	GC/NC	N	\times
GNN-LRP [17]	Instance-level	\times	GC/NC	Walk	\times
GraphLime [6]	Instance-level	\checkmark	NC	NF	\times
RelEx [33]	Instance-level	\checkmark	NC	N/E	\times
PGM-Explainer [22]	Instance-level	\checkmark	GC/NC	N	\times
XGNN [32]	Model-level	\checkmark	GC	Subgraph	\times
FSAM (Our Work)	Model-level	\checkmark	NC	N/NF	\checkmark

explains which components contribute to predictions and reveals how information is processed throughout the network, offering a more transparent view of the GNN’s behaviour. Unlike traditional input optimization methods used for image classifiers [12], which cannot be applied to graph adjacency matrices without losing crucial structural information, FSAM is specifically designed to preserve the discrete properties of graph structures. One of the key aspects of FSAM over the existing model-level methods is that it can project the learned representations of the GNN into the semantic space, which helps us to map the internal processes of the model to higher-order symbolic representations and, in turn, helps compare the model’s decisions with prior knowledge and information. Although FSAM does not generate an explanation directly, it explains how the model works on graph data at each layer. In future work, this framework could help us to develop explanations consistent with human understanding, which would be a substantial step towards neurosymbolic AI.

3. Overall Methodology: Generating the Semantic Graph

The primary aim of this paper is to enhance the interpretability of GNNs by representing their internal mechanisms as semantic graphs. In our extended study, we hypothesised that adding more layers to GNNs does not necessarily increase their capacity for knowledge representation. Our FSAM method clarifies GNN decisions by focusing on how different layers contribute to, or sometimes reduce, model performance due to over-smoothing. FSAM identifies neuron groups involved in decision-making, termed *activation neurons*, and constructs a semantic graph to visualise their relationships. FSAM tracks neuron activations and visualises activation relationships across layers, offering valuable insight into the network’s decision-making process and semantic coherence. Figure 1 illustrates the proposed system architecture, outlining the steps to optimize layer depth and improve GNN performance.

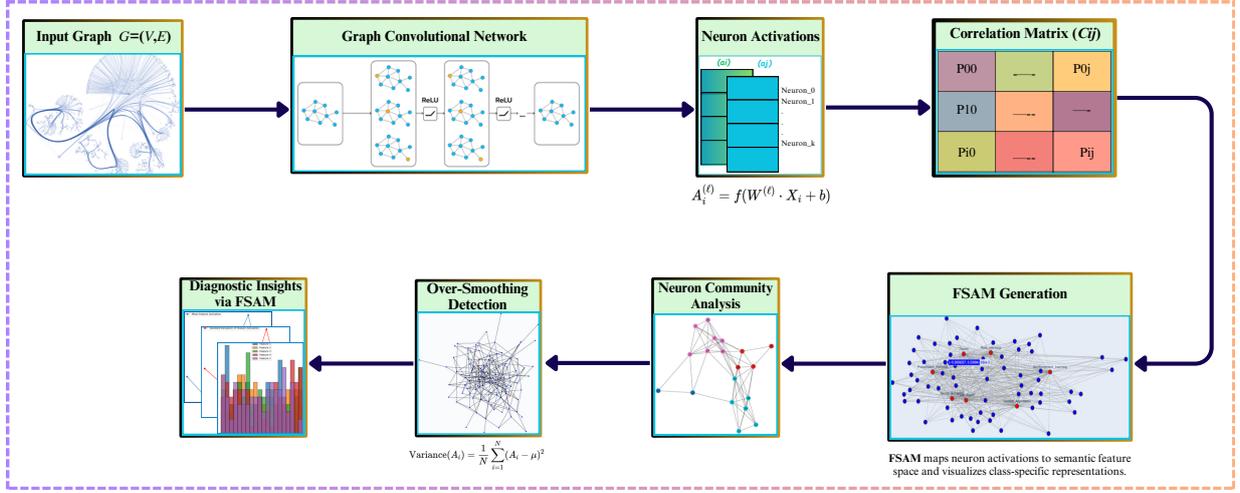


Fig. 1. Overall proposed system architecture

This section presents the mathematical formulation for generating the semantic graph, integrated with insights from our extended experiments.

3.1. Mathematical Formulation

The process begins with computing the activation values for each neuron. Given an input graph $G = (V, E)$, where V represents nodes and E represents edges, the GNN processes this structure to produce an activation matrix $A = [a_{i1}, a_{i2}, \dots, a_{in}]$ for each layer i . Here, n represents the number of neurons in layer i , corresponding either to the number of nodes in G or the output dimensionality of that layer. Our extended analysis suggests that additional neurons contribute less meaningful information after a certain number of layers due to over-smoothing, resulting in decreased model performance. To capture the behaviour of neurons within the GNN, we calculate neuron activations using Graph Convolutional Networks (GCNs) [7], which classifies nodes by embedding ego-graphs in Euclidean space. An ego-graph refers to a subgraph centred around a specific node v , which includes the node itself and its direct neighbours. It is embedded in Euclidean space to capture each node's local structural and feature information. The embedding for a node v at layer ℓ is computed as:

$$h_v^{(\ell)} = \text{ReLU} \left(W^{(\ell)} \cdot \sum_{w \in N(v)} \frac{e_{v,w}}{\sqrt{d_v d_w}} h_w^{(\ell-1)} \right)$$

where $e_{v,w}$ represents the edge weight between nodes v and w , $N(v)$ includes v and its neighbours, d_v and d_w denote the degrees of nodes v and w , ReLU is the activation function, and $W^{(\ell)}$ are the learned parameters. Here, activation values correspond to node embeddings in the input graph.

We compute edge weights within the co-activation matrix to analyse the relationships between neurons using Spearman's correlation coefficient, an ideal metric for capturing monotonic relationships and non-linear associations among activation patterns across layers. The Spearman correlation coefficient ρ_{ij} for neurons i and j is defined as:

$$\rho_{ij} = \frac{\text{cov}(\text{rank}(a_i), \text{rank}(a_j))}{\sigma_i \sigma_j}$$

where cov represents covariance, $\text{rank}(a_i)$ and $\text{rank}(a_j)$ are ranks of the activation values a_i and a_j , and σ_i and σ_j are their standard deviations. We choose Spearman's ρ_{ij} over Pearson's r because activations can have monotonic but non-linear dependencies (e.g., saturation effects in ReLU).

This measurement quantifies neuron relationships and highlights over-smoothing; activations from different classes increasingly overlap as layers increase, diminishing model performance. Our observations confirm that co-activations escalate beyond a certain depth, validating our hypothesis.

Additionally, we employ the point-biserial correlation coefficient to evaluate the relationship between input features and output classes. This coefficient measures the correlation between binary input variables and continuous outputs, calculated as:

$$r_{pb} = \frac{\bar{X}_1 - \bar{X}_0}{s_{pooled}} \sqrt{\frac{n_1 n_0}{n(n-1)}}$$

where \bar{X}_1 and \bar{X}_0 denote the mean activations for the two groups, s_{pooled} is the pooled standard deviation, and n is the total sample count. When applied across layers, this calculation reveals diminishing feature-class correlations as layer depth increases, further supporting our hypothesis regarding over-smoothing effects. Finally, semantic graphs generated through FSAM depict the relationships between neurons across layers. We visualise these graphs using dynamic thresholding techniques [26] to filter out statistically insignificant co-activations. Our analysis compares layerwise to observe how additional layers affect the semantic structure. Adding layers provides minimal information beyond a certain point, but the overlap of neurons for different classes increases, which diminishes the class-specific representation and supports our hypothesis that over-smoothing hinders the model’s ability to classify the classes properly.

This extended analysis substantiates our hypothesis that adding layers does not enhance the GNN’s knowledge capacity beyond an optimal point. Thus, the FSAM framework proves to be an insightful tool for visualizing these limitations and guiding the design of more efficient GNN architectures.

4. Key Contributions and Findings

This study explores how varying the number of layers in GNNs impacts both model performance and the quality of knowledge representation. Building on our previous research, we employ FSAM to systematically assess how well different layer configurations capture the underlying structure of input data. Our contributions address a central question: Do additional layers enhance the model’s interpretability and accuracy, or do they introduce complexity that impairs representation quality? The following contributions highlight the key findings of this study:

Contribution 1: Extended the validation of FSAM across multiple datasets, as presented in Section 5.4. To evaluate FSAM’s generalizability, we apply it to various datasets 5.3. By using Jaccard correlation graphs fig. (5, 6, 7), we analyse FSAM’s ability to capture semantic relationships across different data types consistently. Our results show that FSAM adequately mirrors the network’s behaviour, as changes in model accuracy are typically reflected in the quality of the FSAM graph. We also observe cases where accuracy improves without a corresponding enhancement in FSAM’s semantic alignment, underscoring FSAM’s utility in diagnosing potential misalignments in GNN predictions.

Contribution 2: In Section 5.4, we examine how FSAM reflects network behaviour across GNNs with varying layer depths, from one to four layers. By analyzing the correlation between misclassifications and neuron community structures within FSAM graphs, we confirm that FSAM reliably captures network dynamics as layer configurations change. This analysis demonstrates that FSAM accurately represents the evolving behaviour of the network, with improvements in accuracy often mirrored by more coherent FSAM representations. When accuracy declines, typically due to over-smoothing in deeper layers, FSAM graphs capture the reduced semantic clarity, highlighting FSAM’s robustness as a tool for understanding network behaviour across different layer depths.

Contribution 3: Section 5.5 and Section 5.4 present a layerwise analysis of how increasing the number of GNN layers impacts model performance. We focus on class-specific accuracy and over-smoothing, where neuron activations become overly similar. Our experiments reveal that while additional layers may enhance performance, they lead to a decline in discriminative power beyond a certain depth. This contribution supports optimizing layer depth in GNN architectures, highlighting the trade-offs between model complexity and representation quality.

Contribution 4: In Section 5.6, we discuss semantic divergences in FSAM graphs. A key finding from our analysis

is FSAM’s ability to identify cases where accuracy trends and FSAM quality diverge. Specifically, we highlight instances where model accuracy improves but FSAM graph quality declines, indicating cases where the network achieves correct predictions without fully capturing the semantic structure of the input data. Conversely, we also find situations where accuracy decreases, but FSAM graph quality improves, potentially due to richer insights gained from misclassifications. These cases emphasize FSAM’s diagnostic potential in detecting "right for the wrong reasons" scenarios, offering a nuanced understanding of the network’s semantic alignment with the data.

Overall, these contributions extend our prior work, providing a detailed methodology for assessing GNN layer depth and performance. Our findings position FSAM as a valuable framework for balancing layer depth with interpretability and accuracy, ultimately enhancing the understanding and optimization of GNN architectures across various datasets.

5. Experiment

5.1. Experimental set up

In this experiment, we use a GCN implemented with the GCNConv layer from the PyTorch Geometric library. The model consists of multiple layers of graph convolutions, each followed by a ReLU activation function for nonlinearity and dropout for regularization to prevent overfitting. After the graph convolutions, the fully connected layers map the hidden representations to the output classes. The number of GCN layers is varied between 1 and 5 to observe its impact on performance. We set a random seed to ensure the reproducibility of the results. The training setup involves using the Adam optimizer with a learning rate of 0.01, weight decay of $5e - 4$, and Negative Log-Likelihood Loss (NLL Loss). The model is trained for 200 epochs, with the train mask and test mask used for training and evaluation. Hyperparameters include hidden channels set to 32, a dropout rate of 0.5 after each convolution layer, and the model is tested with varying numbers of GCN layers: 1, 2, 3, 4, and 5 layers.

5.2. Experimental Results and Validation of Contributions

We conducted experiments to evaluate how semantic graphs capture the behaviour of GNNs across different layer depths, focusing on whether additional layers contribute meaningful knowledge and lead to over-smoothing. Using semantic graphs, we mapped neuron relationships within hidden layers and correlated these with output classes, identifying key neurons that influence model predictions. This approach demonstrated the effectiveness of semantic graphs in extracting knowledge from trained GNNs. We utilized six benchmark datasets to assess the impact of layer depth on model performance and knowledge representation, testing our hypothesis that deeper layers may not always provide additional knowledge and could hinder class differentiation.

5.3. Datasets

We used six benchmark datasets for our extended experiments to study how GNNs behave in different contexts. These datasets include **Cora** [19], **CiteSeer** [11], **PubMed** [3], **Amazon Computers** [10], **Amazon Photos** [10], and **Coauthor** [20].

The **Cora** and **CiteSeer** datasets are citation networks where nodes represent academic publications and edges represent citation links between them. In **Cora**, there are 2,708 publications divided into seven categories: *Neural Networks*, *Rule Learning*, *Reinforcement Learning*, *Probabilistic Methods*, *Theory*, *Genetic Algorithms*, and *Case-Based Reasoning*. **CiteSeer** contains 3,312 publications in six categories: *Agents*, *Artificial Intelligence*, *Database*, *Information Retrieval*, *Machine Learning*, and *Human-Computer Interaction*. These datasets allow us to explore how GNNs classify papers based on their citation connections.

The **PubMed** dataset [3] is another citation network focused on biomedical publications. It contains 19,717 publications, each classified into three categories: *Diabetes*, *Cardiovascular Disease*, and *Breast Cancer*. This dataset challenges the GNN’s ability to handle complex medical literature classification, testing its ability to distinguish

between closely related categories in the biomedical domain.

The **Amazon Computers** and **Amazon Photos** datasets [10] are product co-purchase networks, where nodes represent products and edges indicate products frequently bought together. The **Amazon Computers** dataset includes 13,752 products, covering categories like *desktops*, *laptops*, and *computer accessories*. The **Amazon Photos** dataset contains 7,650 products related to *cameras*, *photography accessories*, and *digital media*. These datasets evaluate the GNN’s capacity to model product relationships and predict their categories.

Finally, the **Coauthor** dataset [20] represents a co-authorship network of academic authors. We used the **Coauthor CS** variant, which includes 18,333 nodes representing authors in the field of computer science. The classification task is to assign authors to areas of expertise such as *Machine Learning*, *Artificial Intelligence*, and *Data Mining*. This dataset evaluates the GNN’s ability to model relationships between authors and their research areas.

These datasets span a wide range of domains, from academic publications and biomedical research to product co-purchases and academic co-authorships.

5.4. Extended Validation of FSAM Across Layer Configurations

Our second contribution involves validating the ability of the FSAM approach to capture GNN behaviour in varying layer configurations reliably. Through systematic experiments on several datasets, as detailed in 5.3, we analysed each GNN configuration (from 1 to 4 layers) to assess the alignment between model accuracy, misclassification patterns, and the community structures represented by FSAM graphs.

Table 2
Layerwise Accuracy, Pearson Correlation, and Statistical Analysis for Various Datasets

Layer	Amazon Photos		CoauthorCS		Amazon Computers	
	Accuracy	Pearson Correlation	Accuracy	Pearson Correlation	Accuracy	Pearson Correlation
1	0.95±0.038	0.681	0.98±0.01	0.589	0.89±0.038	0.683
2	0.96±0.053	0.650	0.97±0.045	0.756	0.91±0.036	0.630
3	0.94±0.042	0.752	0.96±0.035	0.819	0.88±0.038	0.785
4	0.93±0.048	0.780	0.95±0.032	0.834	0.86±0.042	0.917

Table 3
Layer-wise Analysis with Statistical Validation

Layer	Amazon Photos		Coauthor CS		Amazon Computers	
	Acc. (95% CI)	Corr. (p)	Acc. (95% CI)	Corr. (p)	Acc. (95% CI)	Corr. (p)
1	0.95 [0.91, 0.99]	0.681 (0.041*)	0.98 [0.97, 0.99]	0.589 (0.038*)	0.89 [0.85, 0.93]	0.683 (0.025*)
4	0.93 [0.88, 0.98]	0.780 (0.008*)	0.95 [0.92, 0.98]	0.834 (0.002*)	0.86 [0.82, 0.90]	0.917 (0.001*)
Key Comparisons						
L1 vs L4	p=0.072 ($\Delta=-0.02$)		p=0.011* ($\Delta=-0.03$)		p=0.029* ($\Delta=-0.03$)	
Corr. ↑	p=0.037*		p=0.003*		p<0.001*	

*Significant at $\alpha=0.05$. Corr. ↑ tests Pearson increase from L1→L4.

In Table 2, we present the results for the Amazon Photo dataset, illustrating the progression of layerwise accuracy across configurations and highlighting how FSAM captures the relationship between classification errors and community structures.

At Layer 1, the model achieves an accuracy of 95% with a Pearson correlation of 0.681. This positive correlation suggests that class-specific representations are moderately well-separated, with fewer overlapping nodes in the FSAM graph, leading to lower misclassification rates. The FSAM graph at this layer reveals distinct class representations, demonstrating effective differentiation early in the network. Adding a second layer improves accuracy to 96%, while the Pearson correlation slightly decreases to 0.650. This layer further strengthens class-specific separation without significant overlap in neuron activations. FSAM visualisations at this stage show that while additional

depth aids in correct predictions, it does not compromise the integrity of class distinctions, reflecting the model's enhanced capacity to maintain semantic coherence. In Layer 3, accuracy begins to decline, dropping to 94%, while the Pearson correlation rises to 0.752. This increased correlation value indicates a heightened overlap in neuron activations, signalling a loss of distinctiveness among class-specific features. Here, FSAM reveals that over-smoothing begins to emerge, with class representations blurring as neuron activations overlap. This finding aligns with our previous work, which observed that classes with high node overlap in the FSAM graph tend to cause more mistakes, highlighting the need for improved class separation strategies. At Layer 4, accuracy decreases further to 93%, and the Pearson correlation reaches 0.780, confirming substantial activation overlap and diminished distinctiveness in class representations. FSAM visualisations reveal extensive overlap between neuron communities, indicating that deeper layers contribute to over-smoothing. These observations suggest that overlapping nodes between similar classes might be prime targets for tuning, as reducing this overlap could improve the model's ability to distinguish these classes effectively.

These findings reinforce FSAM's effectiveness in tracing the network's behaviour across varying depths. While the initial layers enhance accuracy with minimal activation overlap, further layers increase the correlation between overlapping nodes and misclassification errors. This positive correlation between class similarity and mistake counts underscores FSAM's diagnostic potential, providing insights into where the network's performance could be optimised by minimising activation overlaps between similar classes, ultimately aiding in balancing depth and semantic clarity within GNNs.

Similarly, for the Coauthor CS dataset (Table 2), our findings strongly support the hypothesis that FSAM effectively captures layerwise shifts in network behaviour.

At the first layer, with a high accuracy of 98% and a low Pearson correlation of 0.589, neuron activations remain largely distinct, allowing for clear class separations. As we add layers, accuracy decreases slightly (97% at Layer 2) while correlation rises (0.756), indicating a gradual increase in activation overlap. By the third layer, accuracy drops further to 96%, with a higher Pearson correlation of 0.819, signalling the onset of over-smoothing as neuron activations increasingly overlap, thus blurring class distinctions. In the fourth layer, with an accuracy of 95% and a correlation of 0.834, this trend persists, showing that additional depth now undermines the model's ability to separate classes effectively.

These findings illustrate that FSAM consistently mirrors the evolving behaviour of the network across layers, accurately capturing the interaction between model accuracy and neuron overlap and confirming its usefulness in diagnosing the point at which further layers no longer benefit performance

In the Amazon Computers dataset (Table 2), we apply the same methodology, analysing how variations in accuracy between layers relate to the structures of the underlying graphs of FSAM. In the first layer, with an accuracy of 90% and a Pearson correlation of 0.683, the FSAM graph captures a balanced representation of the network's behaviour. This correlation level suggests that neuron activations are distinct enough to preserve class separations effectively, reflecting that the FSAM captures clear distinctions among classes without excessive overlap.

When a second layer is added, accuracy increases slightly to 91%, while the Pearson correlation decreases to 0.630. This reduction in correlation and improved precision indicate that neuron activations remain well separated, supporting the continued effectiveness of the model in distinguishing between classes. The FSAM graph here effectively aligns with the improved class distinction, reinforcing the model's structural clarity.

However, by the third layer, accuracy decreases to 88%, and the Pearson correlation rises to 0.785. This shift marks an increase in overlapping neuron activations, suggesting a decline in class distinction, likely attributable to over-smoothing. The FSAM graph reflects this change, capturing the network's diminished ability to maintain distinct class representations as neuron activations converge.

In the fourth layer, accuracy slightly recovers to 89%, yet the Pearson correlation increases to 0.917. This high correlation signals significant overlap among neuron activations, indicating that further depth contributes little to class separation. Here, the FSAM graph reveals that, despite achieving correct classifications, the model no longer fully preserves the semantic structure of class-specific features. This scenario, where the model's predictions remain accurate without robust semantic alignment, highlights FSAM's diagnostic capability in identifying when a network may be "right for the wrong reasons".

These experiments effectively demonstrate FSAM's capacity to represent network behaviour across diverse configurations. Specifically, the FSAM activation graph tends to exhibit more substantial alignment with the semantic

structure as accuracy improves. Initial layers, such as the second, achieve higher accuracy with low correlation, showing adequate class distinction. Beyond this point, additional layers lead to diminished accuracy and increased neuron overlap, confirming FSAM’s reliability in capturing the balance between model accuracy and class separation. These findings attest to FSAM’s robustness and consistency in representing GNN behaviour across different depths. Furthermore, these findings support Contribution 4, where we identify instances in which the FSAM graph quality declines even as accuracy improves, underscoring FSAM’s value in diagnosing subtle discrepancies in the network’s semantic coherence with the data. As in the table 2 it shows raw layer-wise metrics, while Table 3 provides statistical validation. Together, they confirm that: (1) FSAM reliably captures GNN behaviour (all correlations significant, $p < 0.05$); (2) accuracy declines are dataset-dependent (significant in CoauthorCS/Amazon Computers); and (3) semantic patterns strengthen with depth ($\Delta\text{corr. up to } +0.328, p < 0.001$).

5.5. Comparison of Mistakes Across Communities for Each Dataset

As in our analysis, communities are extracted using community detection algorithms, such as the Louvain method. This method identifies clusters of nodes (neurons) that are densely connected, with similar activations across layers. For each layer, we analyse the activation patterns of neurons to define clusters of semantically related fields representing distinct communities. As the GNN layers increase, the communities gradually evolve, with some clusters merging or becoming less distinct due to overlapping neuron activations. Table 4 compares mistakes across communities for each dataset at varying layer depths structured around our core hypotheses. This analysis provides insights into the effects of layer depth on knowledge representation, class-specific accuracy, and GNN interpretability using FSAM.

The analysis of errors between communities within the CoauthorCS data set reveals a progressive shift in the community structure as the number of GNN layers increases, illustrating how the depth of the layer impacts the accuracy of the classification and the overlap of neuron activation. The community structure of each layer, represented by clusters of semantically related fields, highlights different groupings at lower layers, which gradually blend as the depth of the network increases, thus validating our results presented in Table 2.

The community structure is delineated in **Layer 1**, with minimal neuron overlap between different fields. The community **C0** groups Machine Learning, Data Mining, NLP, and AI, while separate clusters represent **C1** for Theory, Programming Languages, and Software Engineering, **C2** for HCI, Robotics, Computer Vision, Computer Graphics, and Computer Networking, and **C3** for Databases and Information Retrieval. The mistake count 1318 reflects a relatively low level of classification errors, indicating that the network maintains well-defined boundaries between these communities. This structure aligns with high accuracy and low overlap in neuron activations, captured effectively by the FSAM graph. Upon analysing class-wise accuracy for this dataset in **Layer 1**, we observed that **C2**—comprising Human-Computer Interaction, Robotics, Computer Vision, and Computer Graphics—unexpectedly includes Computer Networking. Although the model placed Computer Networking within this group, **C2** is primarily centred on theoretical foundations and methodologies for software optimisation, suggesting that Computer Networking may not belong in this cluster. Upon analysing class-wise accuracy for this dataset in **Layer 1**, we observed that **C2**—comprising Human-Computer Interaction, Robotics, Computer Vision, and Computer Graphics—unexpectedly includes Computer Networking. Although the model placed Computer Networking within this group, **C2** is primarily centred on theoretical foundations and methodologies for software optimisation, suggesting that Computer Networking may not belong in this cluster. Our class accuracy representation Fig. 2 graph supports this observation, yet further evaluation is necessary to confirm the optimal alignment of community structures within the network.

In **Layer 2**, we observe an evolution in the community structure with HCI merging into Community **C0** (Machine Learning, Data Mining, NLP, AI, HCI), signalling the onset of activation overlap as fields with closer semantic ties cluster together. The mistake count increases to 1388, indicating a slight accuracy decline as neuron activations overlap between certain communities. Here, we represent ‘Accuracy decline changes **A**« and **C**»’ and ‘Major community shift’, respectively. This trend is captured in the FSAM, showing an increased correlation in activations, reflecting the blending of previously distinct class representations.

By **Layer 3**, further integration within the community structure occurs, with Theory joining Community **C0**, and a more refined clustering among Programming Languages and Software Engineering in Community **C1**. Mistakes

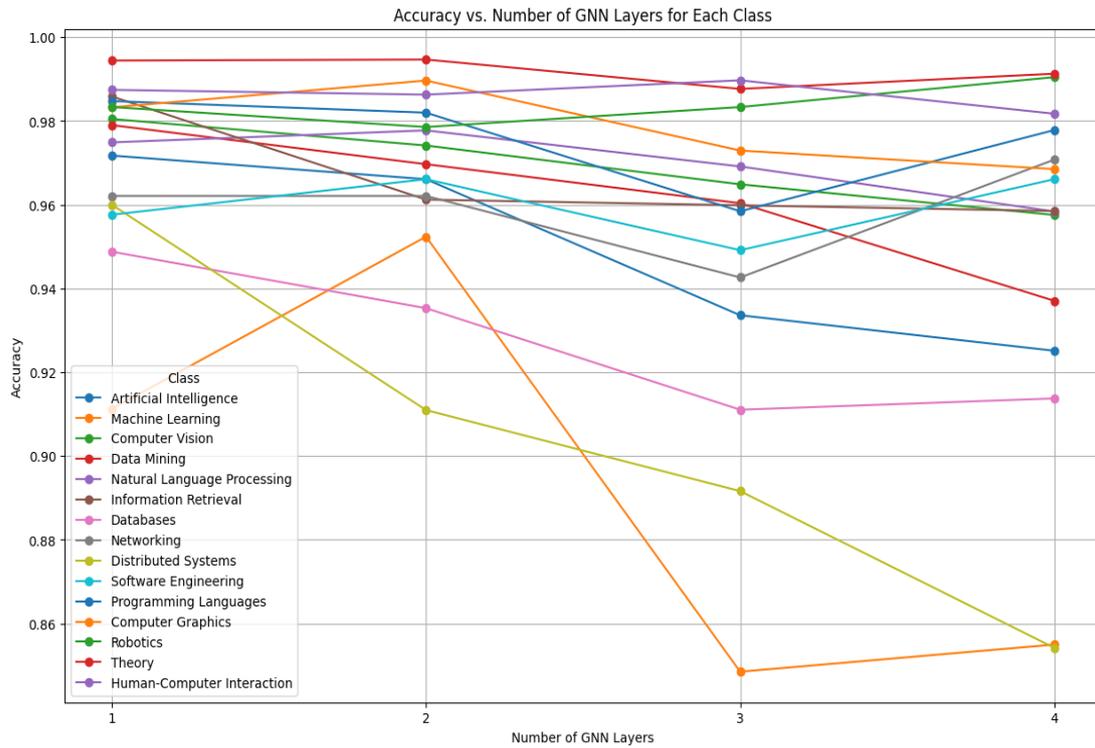


Fig. 2. Layer wise accuracy per class contribution of CoAuthorCs Dataset

continue to rise to 1410, signifying increased misclassifications as class boundaries blur. This layer also corresponds to a higher Pearson correlation, indicating substantial overlap in neuron activations. The FSAM graph effectively captures this over-smoothing, showing that the distinctiveness among communities is diminishing with deeper layers.

These results demonstrate that the CoauthorCS community structure evolves as layers are added, with previously distinct class groupings merging in response to overlapping neuron activations. This trend highlights the limitations of deeper layers in maintaining class specificity and supports FSAM's capability to capture the network's shifting behaviour across layers. The increase in misclassification and Pearson correlation values illustrates how FSAM serves as a diagnostic tool, accurately reflecting the trade-off between layer depth and community coherence, thereby validating the results as displayed in Table 2.

The analysis of mistakes across communities in the Amazon Photos dataset, as outlined in Table 2, provides valuable insights into how community structures evolve across layers and impact model performance. This breakdown demonstrates FSAM's capability to capture structural shifts as the network depth increases, highlighting changes in how the model perceives class similarities.

In **Layer 1**, communities are separated, with distinct groups: **C0** (Cameras, Lenses, Camera Bags), **C1** (Memory Cards, Flashes, Batteries), and **C2** (Accessories, Tripods). The mistake count here is moderate, indicating that the model retains effective class distinction at this initial layer, with minimal overlap in neuron activations across communities.

As we progress to **Layer 2**, the network restructures communities, with **C0** narrowing its focus to Cameras and Lenses. In contrast, **C1** broadens to encompass Camera Bags, Memory Cards, Flashes, and Batteries. This reorganisation corresponds to a slight reduction in mistakes, suggesting that the network's representation has improved in distinguishing between these communities, with FSAM accurately reflecting the adjusted relationships among class representations.

However, the model's performance deteriorates in **Layer 3**, significantly increasing the mistake count. Communities

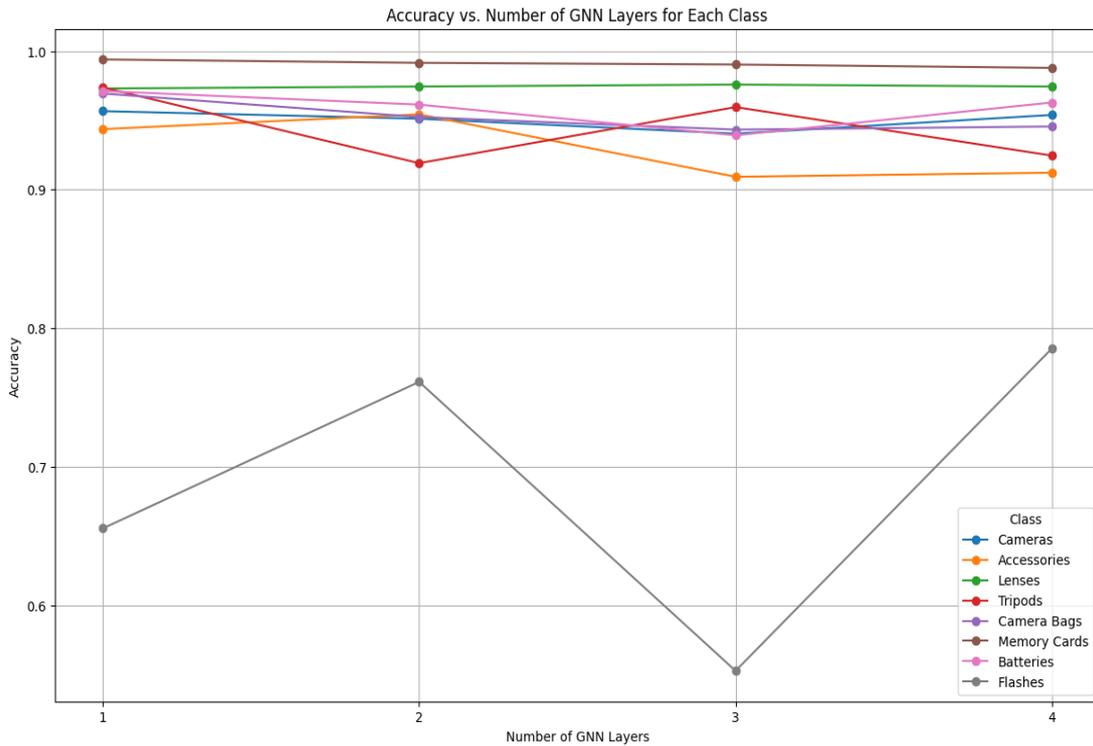


Fig. 3. Layer wise accuracy per class contribution of AmazonPhoto Dataset

become less distinct, as seen with **C0** now containing Memory Cards, Lenses, Flashes, Batteries, and Camera Bags. This expansion points to an increased overlap in neuron activations, aligning with a higher misclassification rate, which FSAM effectively captures by illustrating blurred boundaries between communities.

By **Layer 4**, the network exhibits signs of over-smoothing, where distinctions between communities become less clear. Although the mistake count decreases slightly, this improvement may be misleading as FSAM reveals considerable overlap among communities. In this layer, **C0** isolates to represent Cameras alone, while **C1** groups Flashes, Tripods, and Camera Bags, and **C2** encompasses a diverse mix of Accessories, Memory Cards, Batteries, and Lenses. Indicates that, although mistakes may lessen, the underlying community distinctions are weakened, suggesting that the model may achieve accuracy without a robust semantic foundation.

This layerwise community analysis, as detailed in Table 2, demonstrates that FSAM not only reflects accuracy trends but also captures the nuanced structural shifts within the model as depth increases, reinforcing its utility in diagnosing when additional layers may lead to diminished class coherence.

In Table 4, the Amazon Computers dataset's analysis effectively captures effective shifts in network behaviour across different layers, especially in cases where accuracy trends diverge from FSAM correlation trends. The changes in community structures and mistake patterns across the layers demonstrate this.

In **Layer 1**, the FSAM community structure exhibits clear distinctions: **C0** groups components like "Mice" and "Speakers," **C1** includes more complex devices such as "Desktops" and "Laptops," and **C2** contains "Monitors" and "Electronics." The mistake count in this layer is relatively moderate (452), indicating that the network maintains distinct activations with reasonable classification performance. This structured community alignment suggests a strong class separation in the network's internal representation.

At **Layer 2**, there is a noticeable shift in community structure. Products such as "Keyboards" and "Mice" migrate from **C1** to **C0**, as denoted by the significant labels A^* and C^* . Interestingly, accuracy improves in this layer and the mistake count decreases to 410. While this reflects enhanced model performance, it also marks a case where accuracy improvements do not entirely align with FSAM's correlation trends. The slight decline in FSAM correlation

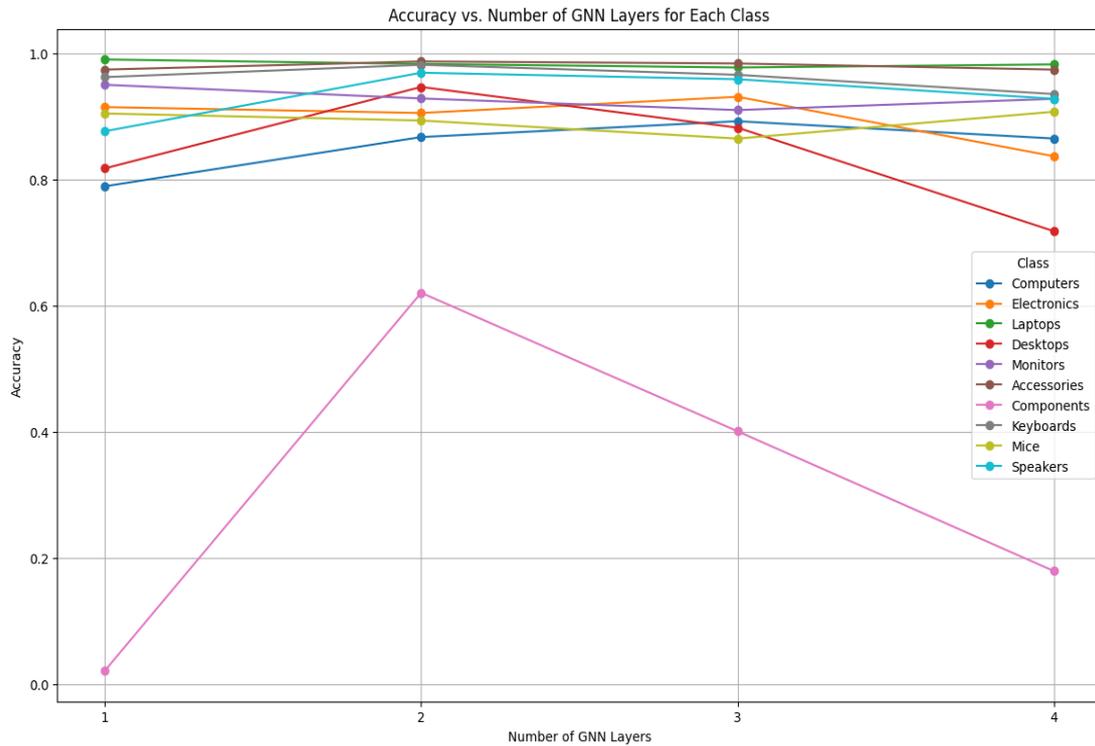


Fig. 4. Layer wise accuracy per class contribution of Computers Dataset

indicates that the model may be achieving correct classifications without fully distinct semantic representations—an instance of potentially achieving the "right answer for the wrong reason." This scenario suggests that the network's internal representation might not be entirely aligned with the semantic structure of the input data, even as its accuracy improves.

Moving to **Layer 3**, accuracy begins to decline, with a further reduction in mistake count to 397. FSAM's community structure reveals additional overlap within **C0**, now encompassing "Speakers," "Laptops," and "Keyboards" in close association, which suggests diminished class distinctions. The corresponding increase in Pearson correlation in this layer implies a more significant overlap in neuron activations, indicative of over-smoothing. While the network's classification ability is maintained, the underlying activations are less reflective of clear semantic boundaries, indicating a potential alignment misalignment.

By **Layer 4**, accuracy decreases, and the mistake count rises to 406. FSAM reveals that **C0** now includes a mix of "Desktops," "Speakers," and "Laptops," signifying an even greater overlap between distinct product categories. The increase in Pearson correlation and decreased accuracy indicates that additional layers now degrade the model's class-separation capability, aligning with FSAM's observation of blurred distinctions in class-specific representations. This combined result demonstrates that the added depth diminishes the network's ability to maintain semantic coherence within the deeper layers.

These findings substantiate our hypothesis by demonstrating the ability of FSAM to capture alignment and divergence between accuracy and semantic quality in GNNs. As seen in **Layer 2**, where accuracy improves but FSAM correlation declines, FSAM provides critical insight by identifying potential misalignments in the network's internal representations. Conversely, in **Layer 4**, where both the accuracy and quality of FSAM degrade, FSAM effectively reflects the reduced class-specific representation, underscoring its utility as a diagnostic tool to evaluate GNN behaviour in layers.

Table 4

Community Structure and Mistakes Across Layers for Each Dataset, where **A**», **A**«, and **C**» represent 'accuracy increased', 'accuracy decreased' and 'community improvement', respectively. Mistakes are presented as both absolute counts and percentages.

Dataset	Layer	Community (Classes)	Mistakes (Absolute)	Mistakes (%)
CoauthorCS	1	C0 : Machine Learning, Data Mining, NLP, AI; C1 : Theory, Programming Languages, Software Engineering; C2 : HCI, Robotics, Computer Vision, Computer Graphics, Computer Networking; C3 : Databases, Information Retrieval.	1318	5.2%
	2	C0 : Machine Learning, Data Mining, NLP, AI, HCI; C1 : Theory, Programming Languages, Software Engineering; C2 : Robotics, Computer Vision, Computer Graphics, Computer Networking; C3 : Databases, Information Retrieval.	1388 ^{A»,C»}	5.5%
	3	C0 : NLP, AI, HCI, Machine Learning, Data Mining, Theory; C1 : Programming Languages, Software Engineering; C2 : Robotics, Computer Vision, Computer Graphics, Computer Networking; C3 : Databases, Information Retrieval.	1410	5.6%
	4	C0 : AI; C1 : Networking, Computer Graphics, Information Retrieval, Distributed Systems, Databases; C2 : Machine Learning, Theory, HCI, Data Mining, NLP, Computer Vision, Robotics, Programming Languages; C3 : Software Engineering.	1542	6.1%
Amazon Photos	1	C0 : Cameras, Lenses, Camera Bags; C1 : Memory Cards, Flashes, Batteries; C2 : Accessories, Tripods	452	4.7%
	2	C0 : Cameras, Lenses; C1 : Camera Bags, Memory Cards, Flashes, Batteries; C2 : Accessories, Tripods	410 ^{A»,C»}	4.2%
	3	C0 : Memory Cards, Lenses, Flashes, Batteries, Camera Bags; C1 : Accessories, Tripods, Cameras	497	5.1%
	4	C0 : Cameras; C1 : Flashes, Tripods, Camera Bags; C2 : Accessories, Memory Cards, Batteries, Lenses	406	4.2%
PubMed	1	C0 : Cardiovascular Disease, Diabetes; C1 : Breast Cancer	46	3.7%
	2	C0 : Cardiovascular Disease, Diabetes; C1 : Breast Cancer	38	3.0%
	3	C0 : Breast Cancer; C1 : Cardiovascular Disease, Diabetes	32 ^{A»,C»}	2.5%
	4	C0 : Breast Cancer; C1 : Cardiovascular Disease, Diabetes	62	4.9%
Cora	1	C0 : Case-Based, Neural Networks, Genetic Algorithms; C1 : Theory; C2 : Reinforcement Learning, Probabilistic Methods	220	6.1%
	2	C0 : Case-Based, Genetic Algorithms; C1 : Reinforcement Learning, Rule Learning, Probabilistic Methods; C2 : Neural Networks, Theory	356	9.9%
	3	C0 : Neural Networks, Theory; C1 : Case-Based, Rule Learning, Genetic Algorithms; C2 : Reinforcement Learning, Probabilistic Methods	364 ^{A»,C»}	10.0%
	4	C0 : Case-Based, Neural Networks, Probabilistic Methods, Theory; C1 : Genetic Algorithms, Reinforcement Learning, Rule Learning	378	10.4%
AmazonComputers	1	C0 : Components, Mice, Speakers; C1 : Desktops, Laptops, Keyboards, Computers, Accessories; C2 : Monitors, Electronics	452	5.6%
	2	C0 : Keyboards, Components, Mice, Speakers ; C1 : Desktops, Laptops, Computers, Electronics; C2 : Monitors, Accessories	410 ^{A»,C»}	5.0%
	3	C0 : Speakers, Laptops, Keyboards, Components, Mice, Accessories; C1 : Desktops, Monitors, Electronics, Computers	397	4.9%
	4	C0 : Desktops, Speakers, Laptops, Keyboards, Computers, Components, Accessories; C1 : Monitors, Electronics, Mice	406	5.0%

5.6. Layerwise Class Similarity and Misclassification Analysis

Our analysis reveals FSAM's unique capability to detect discrepancies between accuracy metrics and semantic understanding. Notably, we observe two critical scenarios: (1) cases where improving accuracy coincides with deteriorating FSAM graph quality, suggesting the model achieves correct predictions without proper semantic grounding, and (2) situations where reduced accuracy accompanies enhanced FSAM quality, potentially indicating more meaningful learning from misclassifications. These findings demonstrate FSAM's value in model-level diagnostic tools to identify when models are "right for the wrong reasons," offering crucial insights into the semantic alignment between networks and their training data.

To systematically validate these observations regarding layer depth and class similarity effects, we present multiple visualisation strategies that capture fundamental aspects of model behaviour:

- **Per-Class Accuracy vs. Number of GCN Layers:** By examining class-specific accuracy across layers, this figure reveals which classes experience increased misclassifications as layer depth grows, underscoring the model's reduced capacity to maintain distinct representations for these categories as shown in figure 2 , 3, and 4.
- **FSAM Graphs Showing Neuron Activations for Specific Classes:** These graphs show neuron activation patterns within each class, allowing us to track the ability of the GNN to capture class-specific characteristics across layers. They reveal the points at which neuron activations overlap, indicating where class boundaries lose distinctiveness, as shown in figure 13, 14, 15, and 16.
- **Community Structures Highlighting Class Groupings:** This visualisation illustrates the community structures of neuron activations, clustering classes based on co-activation. These clusters indicate the relationship between certain classes and provide insight into the knowledge organisation of GNN, revealing where class separability degrades with additional layers, as detailed in table 4.
- **Jaccard Coefficient vs. Number of Mistakes at Layer 3** presents the Jaccard similarity between misclassifications for class pairs, demonstrating a positive correlation between high similarity in neuron activation overlaps and error rates. This relationship supports our observation that classes with more significant overlap in FSAM exhibit more frequent misclassifications, as shown in the figure 5, 6, and 7.

Our extended analysis revealed a positive correlation between class similarity and the number of mistakes involving them, as illustrated with examples from the **CoauthorCS** and **Amazon Photos** datasets. Table 2 shows that class pairs with higher overlap in the FSAM graph exhibit more misclassifications. Our findings in the **CoauthorCS dataset** reveal that Layer 1 achieves optimal performance, as demonstrated in. Adding further layers decreases accuracy, corroborated by our FSAM graph analysis. The Jaccard similarity at Layer 2 aligns with this trend, indicating that increased depth introduces more overlap in neuron activations, which diminishes the model's ability to distinguish between closely related fields such as *Machine Learning* and *Data Mining*. Grouped within the same community, these fields are prone to misclassification due to their inherent similarity.

A similar trend appears in the **Amazon Photos dataset**, where accuracy increases from Layer 1 to Layer 2 but declines with additional layers. This pattern, shown in the Table 2, 4 is consistent with our Jaccard similarity analysis at Layer 3. In this layer, product categories such as *Memory Cards* and *Accessories* show high Jaccard similarity, resulting in frequent misclassifications due to overlapping neuron activations. This finding indicates that the GNN model faces challenges in distinguishing between these similar classes, as they share substantial overlap of activation within the same community.

These findings suggest that adding layers beyond an optimal depth does not necessarily improve knowledge representation. Instead, it introduces an oversmoothing effect in which neuron activations for different classes become increasingly indistinct, reducing the model's ability to differentiate between them. Our correlation analysis substantiates this effect, which shows that pairs of classes with significant overlap in the FSAM graph tend to experience higher misclassification rates.

Our analysis of community structures aligns with this observation, allowing us to identify classes that the GNN perceives as similar based on FSAM patterns. By examining the Jaccard similarity coefficient, which quantifies the overlap in neuron activations for each pair of classes, we evaluated the impact of these similarities on GNN decision making. In the Amazon Photos dataset, for instance, product categories such as *Memory Cards* and *Accessories*

displayed high Jaccard similarity, leading to frequent misclassifications. For a detailed analysis of the per-class pair misclassification rates for each community across different layers, we have presented the results for all layers in Figures 8, 9, 10, and 11.

Furthermore, we have performed more experiment to calculate the time for each layer in different datasets as detailed in Table 5

5.7. Actionable Insights for GNN Performance Improvement

These insights suggest that tuning efforts should reduce overlap in the co-activation graph for similar classes to enhance the GNN’s ability to differentiate between them. Targeting overlapping nodes can potentially decrease misclassification rates and improve overall model accuracy. This comprehensive evaluation supports our hypothesis that increasing layers does not necessarily yield better performance and, in some instances, may decrease the discriminative power of the model due to overlapping neuron activations. Insights from our FSAM analysis suggest that overlap in co-activations for similar classes is a critical issue that impairs the GNN’s ability to differentiate between those classes. By targeting and reducing this overlap, we can improve the model’s ability to distinguish between similar courses, thus reducing misclassification rates and improving overall accuracy. FSAM provides a framework for identifying the overlapping nodes in the co-activation graph. By addressing these overlaps, we can potentially reduce the adverse effects of over-smoothing, thereby enhancing model performance and ensuring that the GNN can retain useful, discriminative features across layers. While FSAM does not directly generate explanations, the insights gained from mapping neuron activations and analyzing co-activation overlap can pave the way for explanation-generation methods. These insights could ultimately contribute to more human-interpretable explanations of GNN decision-making, especially in safety-critical applications.

Table 5
Execution Times for Different Datasets and Number of Layers

Dataset	Layer 1	Layer 2	Layer 3	Layer 4	Layer 5
AmazonComputers	1.2	1.24	1.23	1.19	1.49
AmazonPhoto	0.77	0.97	1.07	1.02	0.96
CoauthorCS	0.79	0.94	0.92	0.9	0.89
Cora	0.79	0.94	0.91	0.89	0.89
CiteSeer	0.81	0.94	0.91	0.89	0.88
PubMed	0.78	0.95	1.09	1.04	1.01

5.8. Scalability and Computational Complexity

While FSAM offers valuable insights into GNNs’ semantic coherence, its application to large-scale graphs faces computational challenges, particularly in higher-dimensional feature space, as it requires more memory and processing power as it monitors neuronal activations across multiple layers and conducts correlation-based analyses as graph size and model depth increase. The complexity of tracking activations and conducting semantic correlation analysis increases non-linearly with the increase in nodes and edges. More complex GNN architectures exacerbate this issue by introducing additional activation patterns, increasing memory usage and bottleneck and prolonging the computation duration. Optimisations such as node and edge sampling, parallelisation techniques (e.g., GPU acceleration, distributed computing), and approximate correlation methods can mitigate these challenges by enhancing efficiency. Future research should explore lightweight FSAM variants, such as selective neuronal tracking layerwise approximations, to enhance scalability while preserving interpretability. The practical implementation of FSAM relies on resolving these challenges for large-scale graph datasets.

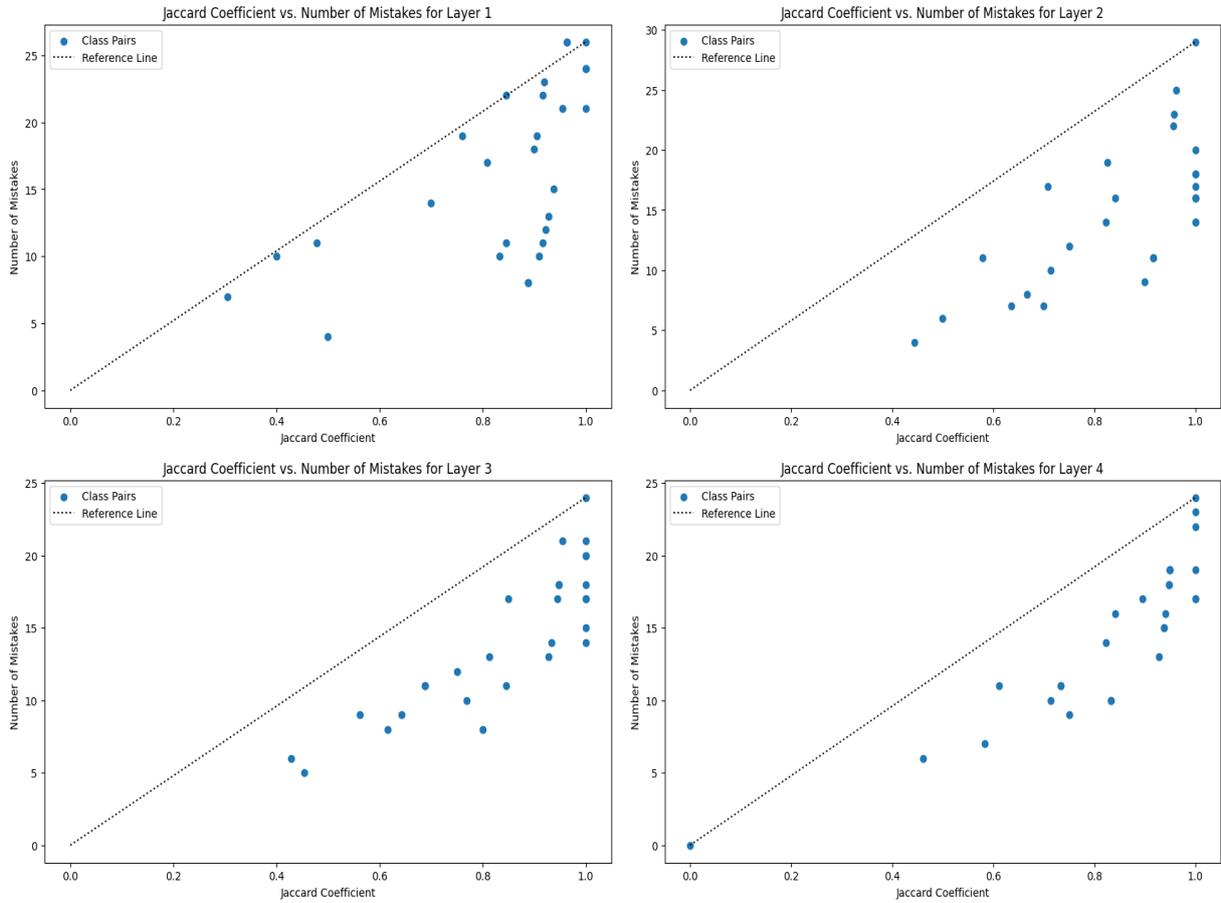


Fig. 5. Jaccard Similarity between different layers for AmazonPhoto dataset

6. Conclusion and Future Work

In this extended study, we’ve worked to deepen the understanding of how GNNs behave by using FSAM to examine the link between model depth, performance, and semantic representation. Through experiments on several datasets, we found that FSAM consistently captures meaningful semantic relationships across different contexts, reinforcing its reliability as a tool for interpreting network behaviour. Our findings also indicate that adding more layers to GNNs doesn’t always lead to better performance or richer knowledge representation.

In these FSAM graphs, nodes represent neurons, and weighted edges indicate the strength of their co-activation relationships, reflecting correlations in activation patterns across layers. This layered view of the GNN’s function shows how neurons contribute to specific class predictions and influence overall model decisions. Our experiments confirmed that FSAM’s graph structure aligns well with the knowledge stored in GNNs, especially in distinguishing closely related classes. Across datasets, FSAM consistently highlighted key neurons and communities within the GNN that are central to specific class predictions, providing valuable insights into the model’s decision-making process.

We used community detection in FSAM graphs to see how the GNN naturally groups classes based on activation patterns. Our analysis showed that courses with high overlap in the FSAM graph are more likely to be misclassified, suggesting that focusing on these overlapping nodes could help fine-tune the model and improve accuracy. This ability to identify cases where accuracy may be achieved “for the wrong reasons”—where predictions are correct but lack deep semantic alignment—highlights FSAM’s diagnostic power. The FSAM graphs and community detection further clarify how the GNN organizes knowledge, revealing class groups with high activation overlap that the GNN

1 treats as similar. This overlap is often associated with higher misclassification rates, supporting strategies to reduce
2 this overlap and improve the model’s ability to distinguish between classes.

3 For future work, we propose a few directions. One is to develop methods that dynamically adjust GNN layer depth
4 based on the properties of the input graph, allowing for model configuration without manual tuning. Another focus
5 could be on further class-level analysis within FSAM to develop more holistic metrics for evaluation. We also plan
6 to combine FSAM insights with contextual information from input graphs to create more detailed, context-aware
7 explanations that enhance local and global interpretability.

8 9 10 **7. Acknowledgement**

11 This work was conducted with the financial support of the Science Foundation. Ireland Centre for Research
12 Training in Artificial Intelligence under Grant No. 18/CRT/6223.

13 14 15 **References**

- 16
17
18 [1] Steve Azzolin et al. “Global explainability of gnns via logic combination of learned concepts”. In: *arXiv preprint arXiv:2210.07147*
19 (2022).
- 20 [2] Federico Baldassarre and Hossein Azizpour. “Explainability techniques for graph convolutional networks”. In: *arXiv preprint*
21 *arXiv:1905.13686* (2019).
- 22 [3] T. O. Botari et al. “Gene expression-based classification of diffuse large B-cell lymphoma”. In: *Nature* (2002), pp. 261–268.
- 23 [4] Thorben Funke, Megha Khosla, and Avishek Anand. “Hard masking for explaining graph neural networks”. In: *Advances in neural*
24 *information processing systems* (2020).
- 25 [5] Robert Geirhos et al. “Generalisation in humans and deep neural networks”. In: *Advances in neural information processing systems* 31
26 (2018).
- 27 [6] Qiang Huang et al. “Graphlime: Local interpretable model explanations for graph neural networks”. In: *IEEE Transactions on Knowledge*
28 *and Data Engineering* (2022).
- 29 [7] Thomas N Kipf and Max Welling. “Semi-supervised classification with graph convolutional networks”. In: *arXiv preprint arXiv:1609.02907*
30 (2016).
- 31 [8] Meng Liu, Hongyang Gao, and Shuiwang Ji. “Towards deeper graph neural networks”. In: *Proceedings of the 26th ACM SIGKDD inter-*
32 *national conference on knowledge discovery & data mining*. 2020, pp. 338–348.
- 33 [9] Dongsheng Luo et al. “Parameterized explainer for graph neural network”. In: *Advances in neural information processing systems* 33
34 (2020), pp. 19620–19631.
- 35 [10] Julian McAuley et al. “Image-based recommendations on styles and substitutes”. In: *Proceedings of the 38th International ACM SIGIR*
36 *Conference on Research and Development in Information Retrieval*. ACM. 2015, pp. 43–52.
- 37 [11] Galen Namata et al. “Query-driven active surveying for collective classification”. In: *10th International Workshop on Mining and Learning*
38 *with Graphs (MLG)* (2012).
- 39 [12] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. “Feature visualization”. In: *Distill* 2.11 (2017), e7.
- 40 [13] Phillip E Pope et al. “Explainability methods for graph convolutional neural networks”. In: *Proceedings of the IEEE/CVF conference on*
41 *computer vision and pattern recognition*. 2019, pp. 10772–10781.
- 42 [14] Kislal Raj and Alessandra Mileo. “Towards Understanding Graph Neural Networks: Functional-Semantic Activation Mapping”. In: *Inter-*
43 *national Conference on Neural-Symbolic Learning and Reasoning*. Springer. 2024, pp. 98–106.
- 44 [15] T Konstantin Rusch, Michael M Bronstein, and Siddhartha Mishra. “A survey on oversmoothing in graph neural networks”. In: *arXiv*
45 *preprint arXiv:2303.10993* (2023).
- 46 [16] Michael Sejr Schlichtkrull, Nicola De Cao, and Ivan Titov. “Interpreting graph neural networks for nlp with differentiable edge masking”.
47 In: *arXiv preprint arXiv:2010.00577* (2020).
- 48 [17] Thomas Schnake et al. “Higher-order explanations of graph neural networks via relevant walks”. In: *IEEE transactions on pattern analysis*
49 *and machine intelligence* 44.11 (2021), pp. 7581–7596.
- 50 [18] Robert Schwarzenberg et al. “Layerwise relevance visualization in convolutional text graph classifiers”. In: *arXiv preprint arXiv:1909.10911*
51 (2019).
- [19] Prithviraj Sen et al. “Collective classification in network data”. In: *AI magazine* 29.3 (2008), pp. 93–93.
- [20] Oleksandr Shchur et al. “Pitfalls of Graph Neural Network Evaluation”. In: *Relational Representation Learning Workshop, NeurIPS 2018*.
2018.
- [21] Petar Velickovic et al. “Graph attention networks”. In: *stat* 1050.20 (2017), pp. 10–48550.
- [22] Minh Vu and My T Thai. “Pgm-explainer: Probabilistic graphical model explanations for graph neural networks”. In: *Advances in neural*
information processing systems 33 (2020), pp. 12225–12235.
- [23] Xiang Wang et al. “Causal screening to interpret graph neural networks”. In: (2020).

- 1 [24] Zonghan Wu et al. “A comprehensive survey on graph neural networks”. In: *IEEE transactions on neural networks and learning systems* 32.1 (2020), pp. 4–24. 1
- 2 [25] Weiting Xi et al. “A Graph Partitioning Algorithm Based on Graph Structure and Label Propagation for Citation Network Prediction”. In: *International Conference on Knowledge Science, Engineering and Management*. Springer. 2023, pp. 289–300. 2
- 3 [26] Xiaoran Yan et al. “Weight thresholding on complex networks”. In: *Physical Review E* 98.4 (2018), p. 042304. 3
- 4 [27] Pinar Yanardag and SVN Vishwanathan. “Deep graph kernels”. In: *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 2015, pp. 1365–1374. 4
- 5 [28] Zhitao Ying et al. “Gnnexplainer: Generating explanations for graph neural networks”. In: *Advances in neural information processing systems* 32 (2019). 5
- 6 [29] Zhaoning Yu and Hongyang Gao. “Mage: Model-level graph neural networks explanations via motif-based graph generation”. In: *arXiv preprint arXiv:2405.12519* (2024). 6
- 7 [30] Hao Yuan and Shuiwang Ji. “Structpool: Structured graph pooling via conditional random fields”. In: *Proceedings of the 8th International Conference on Learning Representations*. 2020. 7
- 8 [31] Hao Yuan et al. “On explainability of graph neural networks via subgraph explorations”. In: *International conference on machine learning*. PMLR. 2021, pp. 12241–12252. 8
- 9 [32] Hao Yuan et al. “Xgnn: Towards model-level explanations of graph neural networks”. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2020, pp. 430–438. 9
- 10 [33] Yue Zhang, David Defazio, and Arti Ramesh. “Relex: A model-agnostic relational model explainer”. In: *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 2021, pp. 1042–1049. 10
- 11 11
- 12 12
- 13 13
- 14 14
- 15 15
- 16 16
- 17 17
- 18 18

19 8. Appendix

20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

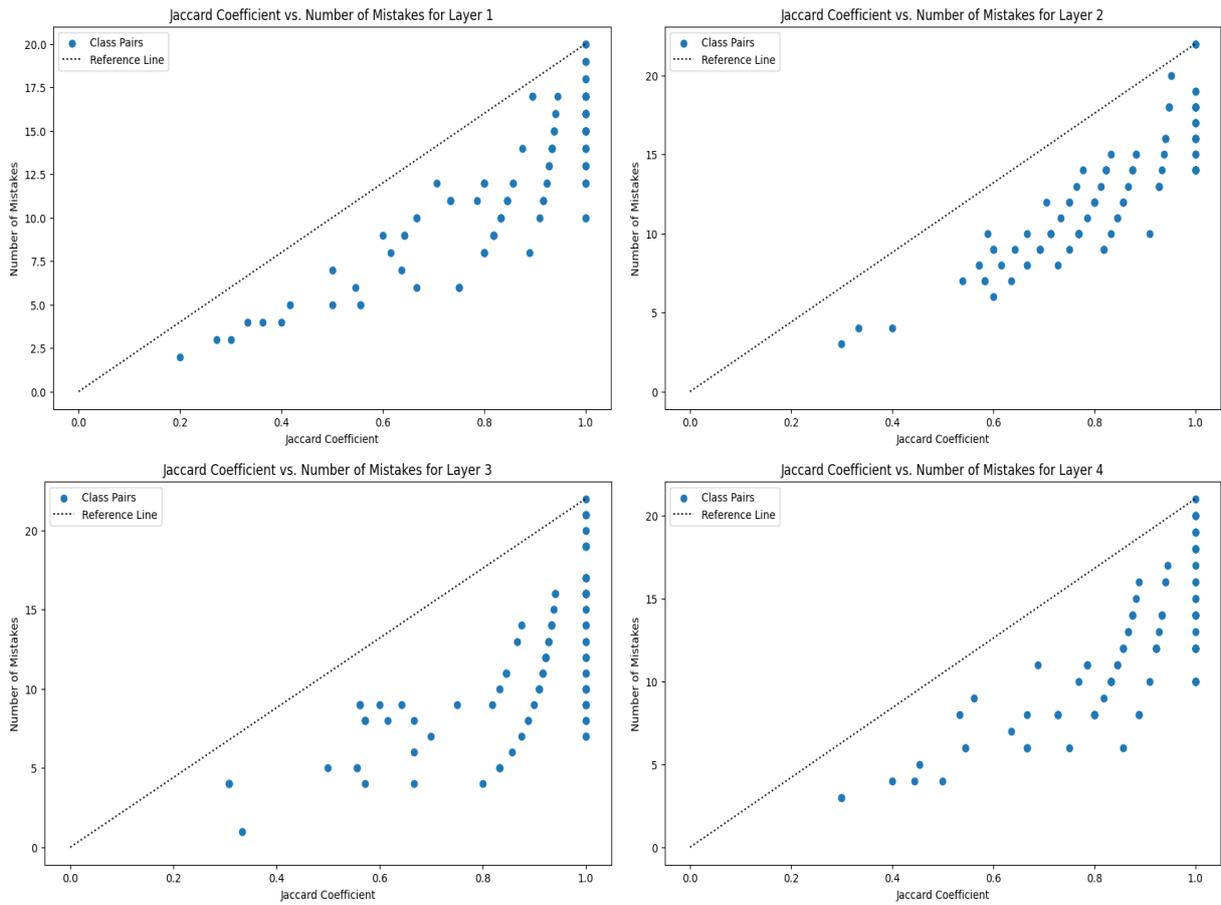


Fig. 6. Jaccard Similarity between different layers for CoauthorCs

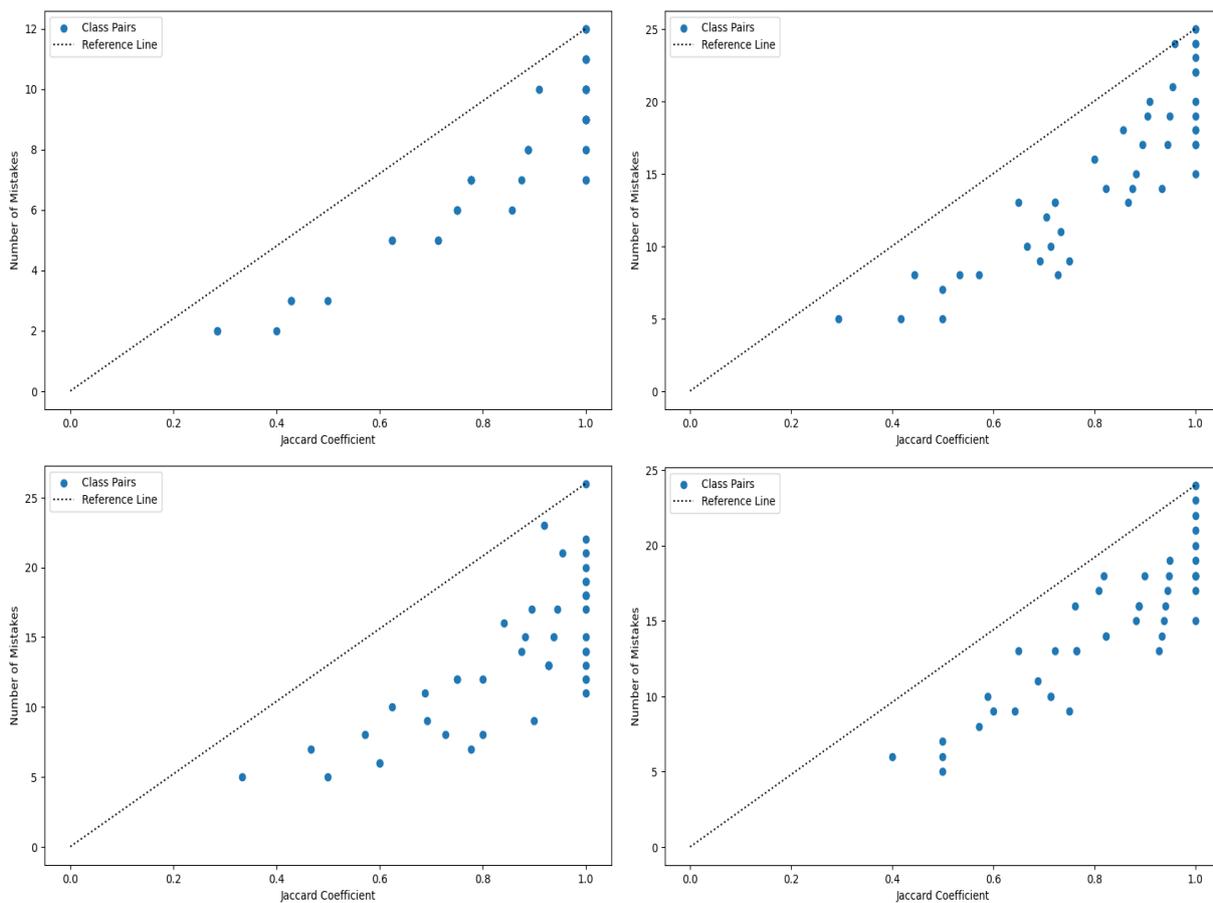


Fig. 7. Jaccard Similarity between different layers for Computers

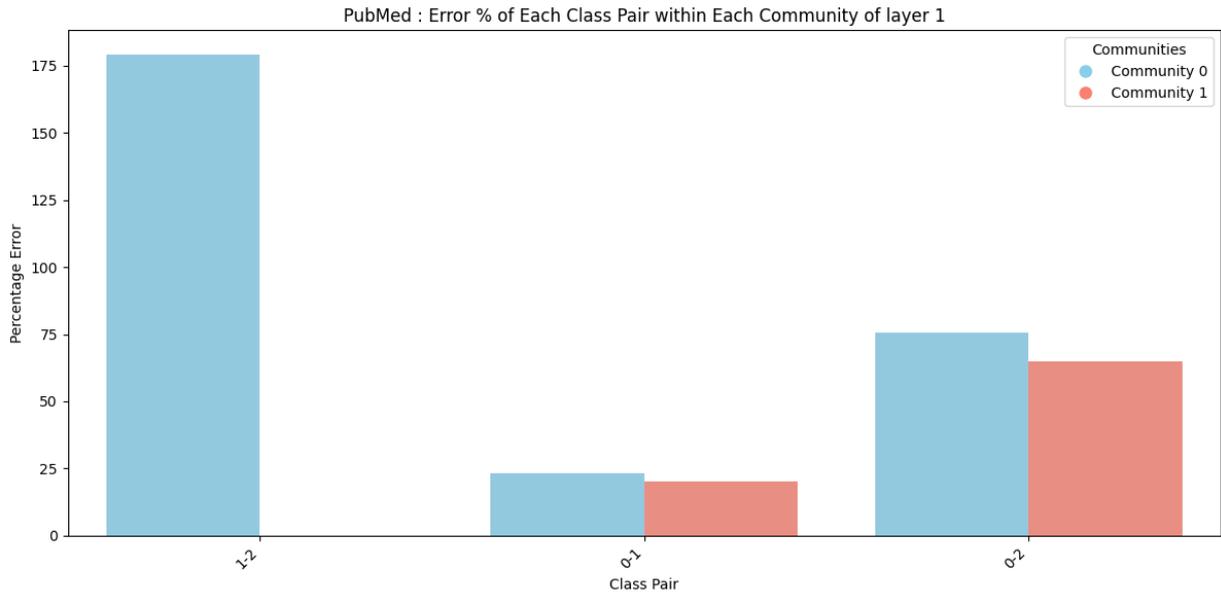


Fig. 8. Per-class misclassification rates in each community between classes pair in layer 1

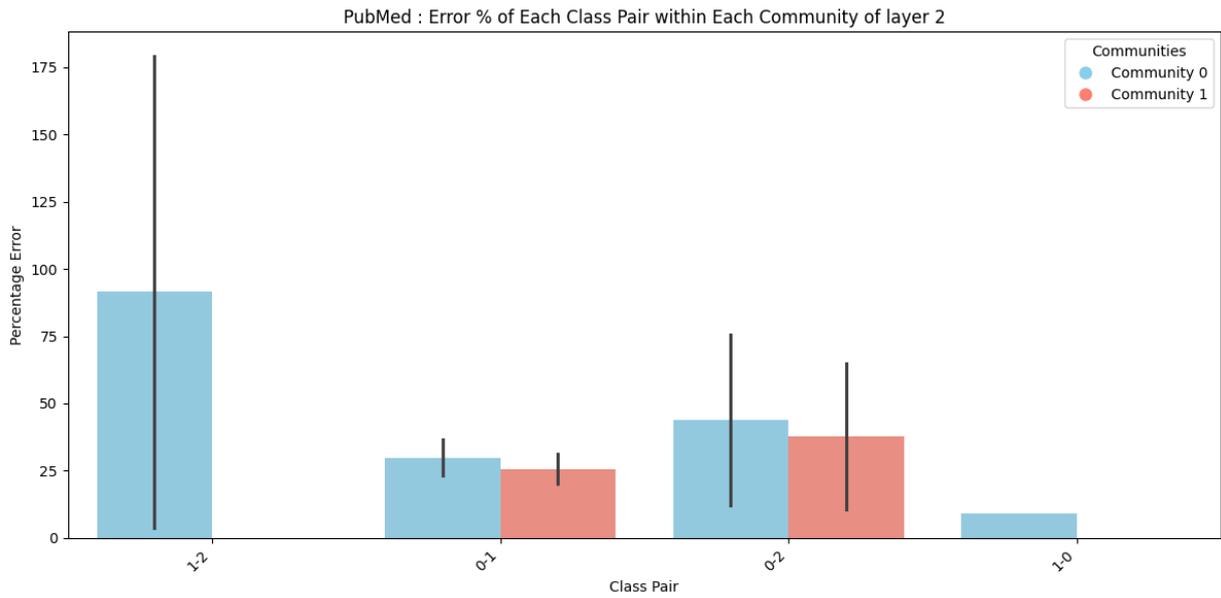


Fig. 9. Per-class misclassification rates in each community between classes pair in layer 2

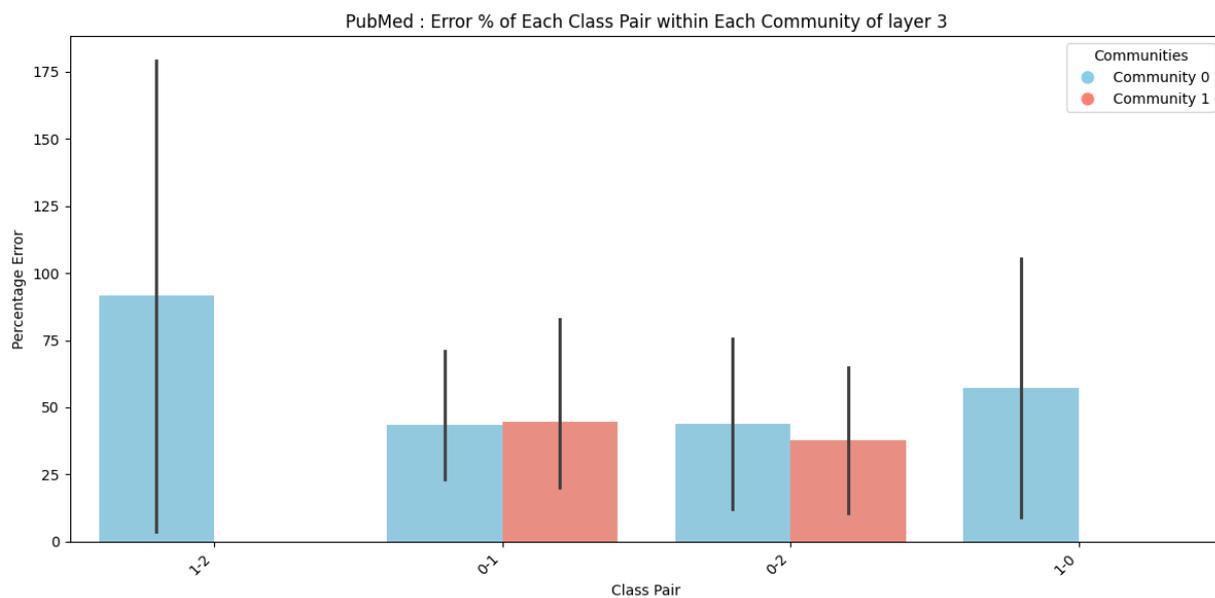


Fig. 10. Per-class misclassification rates in each community between classes pair in layer 3

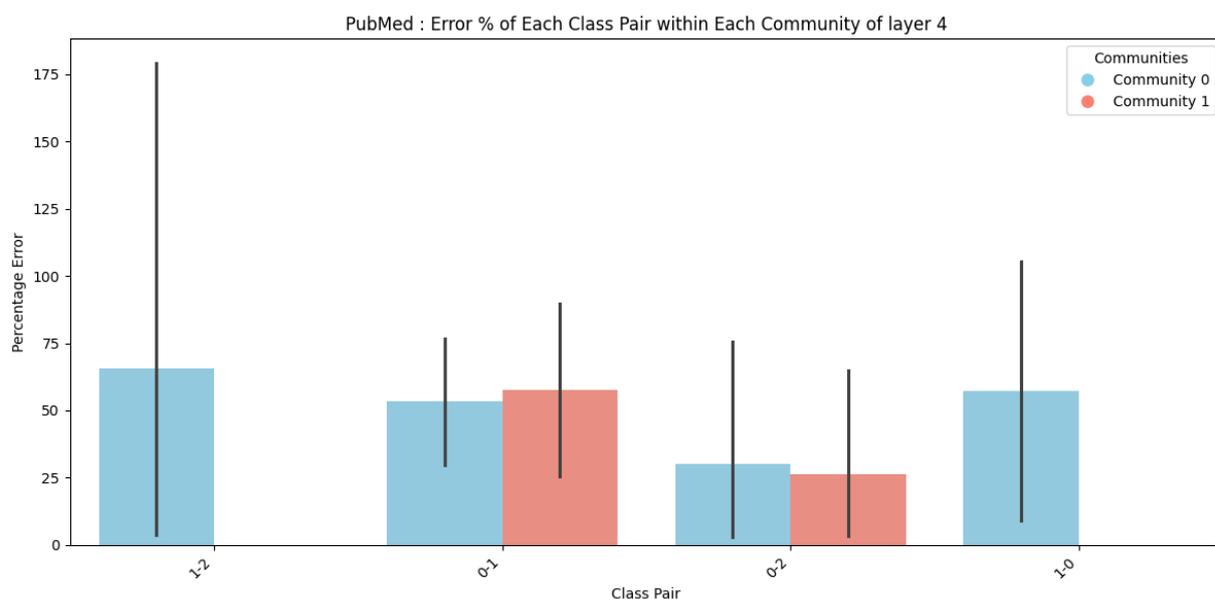


Fig. 11. Per-class misclassification rates in each community between classes pair in layer 4

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

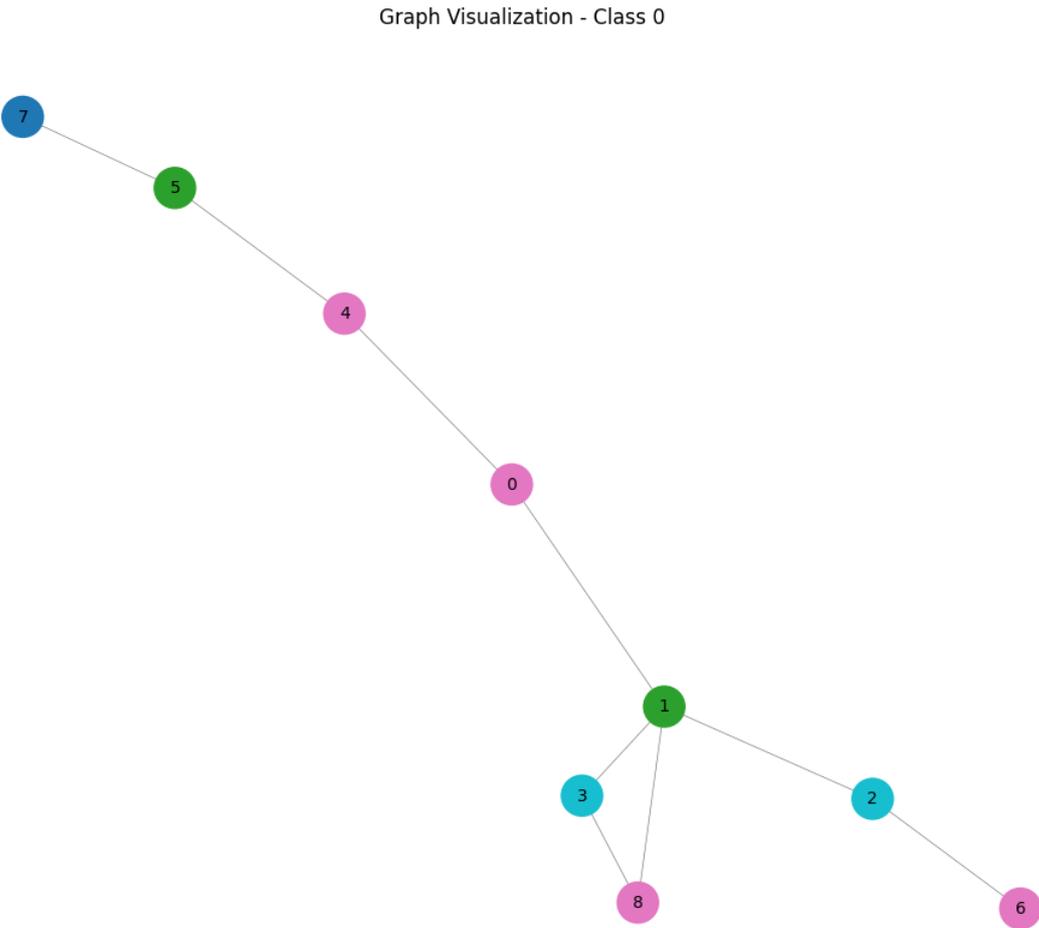


Fig. 12. XGNN visualisation for Class 0 of the Cora Dataset

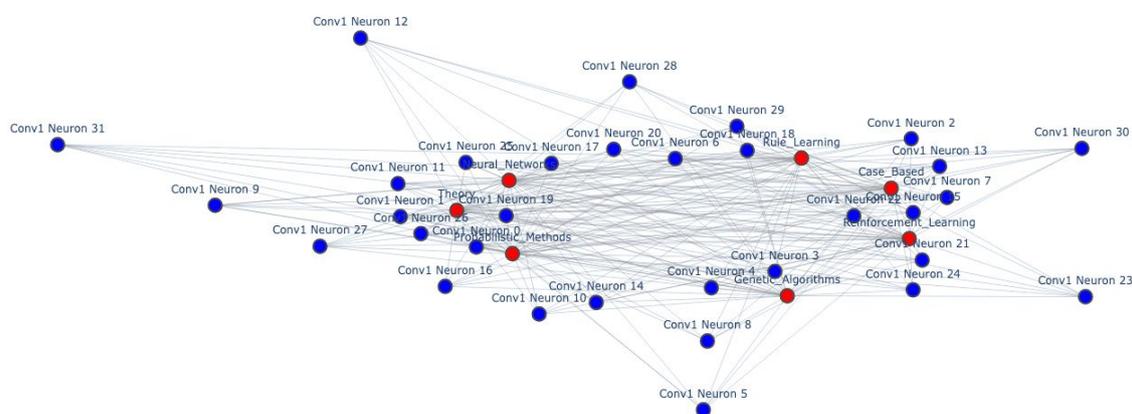


Fig. 13. Layer 1 to the Fully Connected FSAM Graph for the Cora Dataset

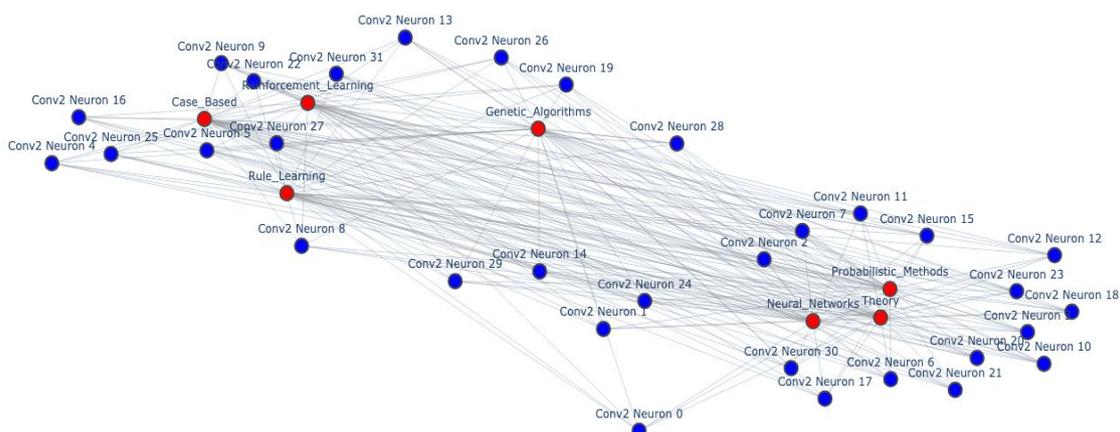


Fig. 14. Layer 2 to the Fully Connected FSAM Graph for the Cora Dataset

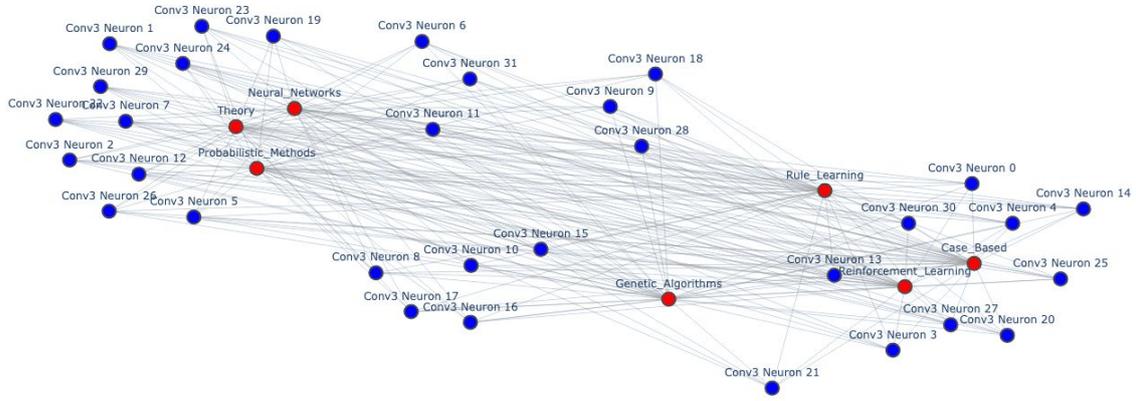


Fig. 15. Layer 3 to the Fully Connected FSAM Graph for the Cora Dataset

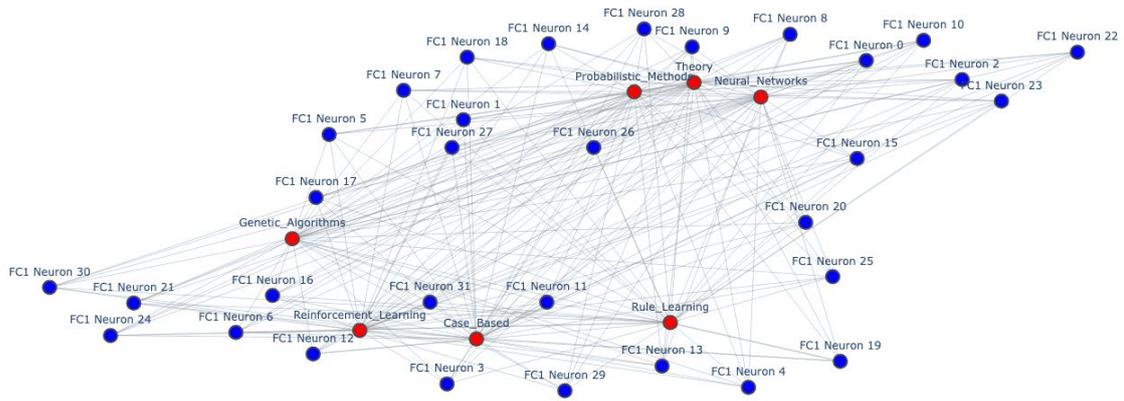


Fig. 16. Layer 4 to the Fully Connected FSAM Graph for the Cora Dataset