

Global Concept-Based Interpretability for Graph Neural Networks via Neuron Analysis

Han Xuanyuan, Pietro Barbiero, Dobrik Georgiev, Lucie Charlotte Magister, Pietro Liò

University of Cambridge
hx263@cantab.ac.uk {pb737, dgg30, lcm67, pl219}@cam.ac.uk

Abstract

Graph neural networks (GNNs) are highly effective on a variety of graph-related tasks; however, they lack interpretability and transparency. Current explainability approaches are typically local and treat GNNs as black-boxes. They do not look inside the model, inhibiting human trust in the model and explanations. Motivated by the ability of neurons to detect high-level semantic concepts in vision models, we perform a novel analysis on the behaviour of individual GNN neurons to answer questions about GNN interpretability. We propose a novel approach for producing global explanations for GNNs using neuron-level concepts to enable practitioners to have a high-level view of the model. Specifically, (i) to the best of our knowledge, this is the first work which shows that GNN neurons act as concept detectors and have strong alignment with concepts formulated as logical compositions of node degree and neighbourhood properties; (ii) we quantitatively assess the importance of detected concepts, and identify a trade-off between training duration and neuron-level interpretability; (iii) we demonstrate that our global explainability approach has advantages over the current state-of-the-art – we can disentangle the explanation into individual interpretable concepts backed by logical descriptions, which reduces potential for bias and improves user-friendliness.

Introduction

Graph neural networks (GNNs) have shown promising results in a variety of tasks where graph-structured data is prevalent (Wieder et al. 2020; Sanchez-Gonzalez et al. 2020; Wu et al. 2021). Like most neural methods, GNNs lack explainability and interpretability. Their complex decision-making is effectively a black-box process to a human observer (Dai et al. 2022), which inhibits human trust in critical applications such as healthcare (Amann et al. 2020; Bussmann et al. 2021; Deeks 2019). Although there has been significant research on the behaviour of individual neurons of convolutional neural networks (CNNs) and recurrent neural networks (RNNs) (Bau et al. 2017; Dalvi et al. 2019; Samek et al. 2019), existing approaches for explaining GNNs tend to treat the model as a black-box and only consider input feature attribution (Dai et al. 2022), ignoring the potential of concepts as a way of interpreting models, which have been

shown to better align with human intuition due to the ability to present an explanation in terms of interpretable units (Kim et al. 2018; Koh et al. 2020; Ghorbani et al. 2019). We perform an investigation into the presence of concepts in GNNs, and propose a new approach for global GNN explainability. We make the following contributions:

1. We show that GNN neurons have alignment with concepts formulated as logical compositions of node and neighbourhood properties, and act as concept detectors which detect *interpretable units of information* such as chemical substructures and social network motifs.
2. We propose new evaluation metrics adapted for GNN neurons. We quantitatively analyse the relationship between neuron importance and interpretability, and also show that there is a trade-off between model training and interpretability: GNNs trained for longer become harder to interpret, even if model performance plateaus.
3. We compare our global GNN explanations against the existing state-of-the-art approach and show that our method possesses advantages for GNN users, including being concept-based, providing explanations backed by logical descriptions, and being inherently noise-robust.

Background

Explainability for GNNs is a relatively new direction. Most approaches focus on local post-hoc explainability using input attribution via methods such as mutual information maximisation. For example, Ying et al. (2019) proposed GNNExplainer. Many other works focus on providing local explanations which are more robust and faithful (Dai et al. 2022; Luo et al. 2020; Vu and Thai 2020). On the other hand, a global or *model-level* explanation tells the practitioner how the model behaves as a whole. Model-level (or global) approaches are less explored for GNNs. The current state-of-the-art model-level approach is XGNN, which uses reinforcement learning to generate a graph that maximises the prediction of a certain class (Yuan et al. 2020; Dai et al. 2022). Unfortunately, since their form of explanation is a graph generated using a black-box neural network, the meaning of the graph is not always intelligible and leaves interpretation to the user, which can lead to human bias (Himmelhuber et al. 2021). However, a recent trend in explainability is to use *concepts*, which have been shown to

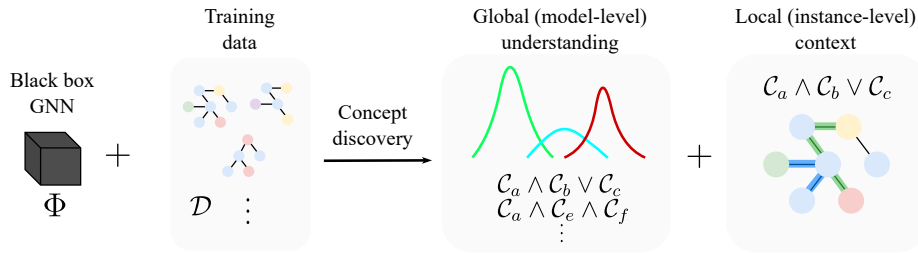


Figure 1: The proposed framework. Given a black-box GNN Φ and training set of graphs \mathcal{D} , we (i) discover interpretable neurons in Φ that act as concept detectors by performing a search over \mathcal{D} ; (ii) utilise the discovered concepts to produce global explanations; (iii) adapt the concepts to specific local instances from \mathcal{D} to produce visualisations to aid the practitioner.

make explanations more plausible and less dependent on human interpretation due to their ability to express a model’s decision making in terms of high-level interpretable units of information (Kim et al. 2018; Koh et al. 2020; Ghorbani et al. 2019). We assume the definition of concepts proposed by Ghorbani et al. (2019), i.e. a concept should ideally satisfy the properties of *meaningfulness*, *coherency* and *importance*. Concepts have been applied to GNNs by Magister et al. (2021), who treat a concept as a k -means cluster in the activation space, and also to the task of algorithmic reasoning via a concept bottleneck layer (Koh et al. 2020; Georgiev et al. 2022). Moreover, in the vision and language domains much work on probing individual model neurons has been done (Bau et al. 2017; Dalvi et al. 2019; Mu and Andreas 2020), and it has been found that individual neurons can behave as *concept detectors*. The potential of such an approach to address GNN interpretability has been unexplored.

Graph Neural Networks

We define a graph $G = (E, V)$ as a set of edges E and nodes V . A GNN model is a function $\Phi(\mathbf{X}, \mathbf{A})$ such that $\mathbf{X} \in \mathbb{R}^{D \times |V|}$ is a node feature matrix, where each column \mathbf{x}_i is the D -dimensional feature vector of node i , and $\mathbf{A} \in \{0, 1\}^{|V| \times |V|}$ is an adjacency matrix which represents the edges E . Although many GNN architectures exist, most can be categorised under *message passing* GNNs (Veličković 2022). At each layer of message passing, node representations $\mathbf{H} \in \mathbb{R}^{D \times |V|}$ are updated using message passing function f , i.e. $\mathbf{H}^{(l+1)} = f(\mathbf{H}^{(l)}, \mathbf{A})$. We present our method assuming that the GNN follows this general framework. Furthermore, although there exists many ways of initialising the node representations, we assume that $\mathbf{H}^{(0)} = \mathbf{X}$. We shall refer to $\mathbf{H}^{(l)}$ as the activations at the l -th layer, and $\mathbf{H}^{(l)}[i, :]$ as the i -th *neuron* at that layer¹.

Methodology

In this section we outline the proposed method. A visual overview is shown in Figure 1.

¹We use slicing notation. $\mathbf{H}[i, :]$ indicates the i -th row of \mathbf{H} , and $\mathbf{H}[:, j]$ the j -th column.

Concept Discovery

With inspiration from the concept extraction method by Mu and Andreas (2020) designed for vision and language models, we propose to perform a compositional concept search over the activation space of a GNN. Formally, letting \mathcal{G} be the space of input graphs, we treat concepts as functions $C : \mathcal{G} \rightarrow \{0, 1\}^{|V|}$ which outputs a boolean value indicating whether each node is part of the concept, i.e. binary concept *masks*. We define a concept C to be *interpretable* if there exists a natural language description of C such that a practitioner is able to accurately predict the output of C given any graph $G \in \mathcal{G}$ from solely knowing this description. Hence, concepts such as *nodes that have exactly one neighbour* would be interpretable. Furthermore, we refer to the *space* of concepts as a set \mathcal{C} which contains all concepts that are interpretable.

Our goal is to determine whether neurons are concept detectors. For a graph G , let V_G be the nodes in G , $C(G)[i]$ indicate the mask value for the i -th node, and $\mathbf{H}_G^{(l)}$ be the activations of the model taking G as input. Note that $\mathbf{H}_G^{(l)}[k, i]$ indicates the k -th neuron activation of the i -th node. We formally define a neuron k to be a concept detector if

$$\exists C. \forall G \in \mathcal{G}. \left[\forall i \in V_G. \left[C(G)[i] \approx \tau_k \left(\mathbf{H}_G^{(l)}[k, i] \right) \right] \right]. \quad (1)$$

where $\tau_k : \mathbb{R}^m \rightarrow [0, 1]^m$ is an element-wise thresholding function specific for neuron k . This formulation assumes that there are concepts with near-perfect alignment with the activations, which in practice is unrealistic. Therefore, we relax the definition and consider whether there are concepts that approximately resemble the behaviour of a neuron k :

$$\exists C. \mathbb{E}_{G \sim \mathcal{G}} \left[\text{Div} \left(C(G), \mathbf{H}_G^{(l)}[k, :, \tau_k] \right) \right] \approx 0. \quad (2)$$

Here, Div is some measure of divergence between the concept mask $C(G)$ and the activations thresholded using τ_k . Following previous works on image models (Mu and Andreas 2020; Bau et al. 2017) that have used intersection over union (IOU) scores, we choose a similar metric

$$\text{Div}(\mathbf{a}, \mathbf{b}, \tau) = - \frac{\sum_i (\mathbf{a} \cap \tau(\mathbf{b}))_i}{\sum_i (\mathbf{a} \cup \tau(\mathbf{b}))_i} \cdot \frac{\mathbf{a} \cdot \tau(\mathbf{b})}{\sum_i b_i}, \quad (3)$$

which is the IOU score scaled by a factor representing how much of the activation signal is incorporated into the concept. Note that $\mathbf{a} \cdot \tau(\mathbf{b})$ represents the total activations within

the concept mask, and $\sum b_i$ is the sum of all activations in the graph. Following previous work (Mu and Andreas 2020; Bau et al. 2017), we apply dynamic thresholding to obtain τ . The search objective becomes the process of finding a concept that maximises a score $f(C, k)$ for neuron k .

$$\arg \max_C f(C, k) \text{ s.t. } f(C, k) = \min_{\tau_k} |\mathcal{D}|^{-1} \sum_{G \in \mathcal{D}} \text{Div} \left(C(G), \mathbf{H}_G^{(L)}[k, :], \tau_k \right) \quad (4)$$

in which we search through the concept space \mathcal{C} . This is tractable if $|\mathcal{C}|$ is small. However, following Mu and Andreas (2020) we desire to search for *compositional* concepts, i.e. those which are a logical composition over base concepts. Formally, we define a compositional concept of length k as an expression of the form

$$C_1 \oplus_1 C_2 \oplus_2 \dots \oplus_3 C_k, \quad (5)$$

where $\{C_i\}_{i=1:k} \subseteq \mathcal{C} \cup \{-C \mid C \in \mathcal{C}\}$ and $\{\oplus_i\}_{i=1:k} \subseteq \{\wedge, \vee\}$, which is intractable for large k . A beam search is performed over the space of compositional concepts using the concept divergence as a priority metric to prune the less desirable concepts. The full algorithm and implementation details are provided in Appendix A². Note that in practice, we use a GPU-based implementation to perform this search on all neurons in parallel.

Concept-Based Model-Level Explanations

Given GNN Φ and target class y , we propose to provide model-level explanations using the concept-extraction framework. We follow the concept-bottleneck paradigm (Koh et al. 2020; Barbiero et al. 2022) of separating the top-level of Φ (the section after the final message passing round) from the rest of the model. Denoting this top level with Φ_T , we propose to produce an explanation for Φ_T as a set of neurons S_{glob} that globally explains the behaviour of Φ for class y . Then, we align the explanation with concepts by finding $\mathcal{E}_{\text{glob}} = \{\arg \max_{C \in \mathcal{C}} f(k, C) \mid k \in S_{\text{glob}}\}$, which is the set of highest scoring concepts for neurons in S_{glob} . For models where Φ_T is a single perceptron, we propose to select the neurons with positive weight contributions for class y .

Due to the importance of visualisation for interpretability (Samek et al. 2019), for each concept in $\mathcal{E}_{\text{glob}}$ we also produce a visualisation to accompany the concept using *activation maps* which are a common approach for vision models (Selvaraju et al. 2019; Zhou et al. 2016). However, a clear distinction is that our activation maps are concept-based and operate over graphs (rather than images) – we therefore refer to them as *graph concept activation maps*.

Definition (Graph concept activation map). Given GNN Φ with L layers, graph $G = (V, E)$, target explanation class y and neuron k , the concept activation map for (Φ, G, y, k) is $(C_k, \eta_V^k, \eta_E^k)$ where C_k is the best-matching concept for neuron k . Also, $\eta_V^k \in \mathbb{R}^{|V|}$ and $\eta_E^k \in \mathbb{R}^{|V| \times |V|}$ are importance masks over the nodes and edges respectively. Inspired

²The full paper along with the appendix can be found on ArXiv: <https://arxiv.org/abs/2208.10609>. The code is available at <https://github.com/xuyhan/gnn-dissect>.

by activation maps on CNNs (Selvaraju et al. 2019; Zhou et al. 2016), we take $\eta_V = \frac{\partial \Phi(G)_y}{\partial n_G^k} \cdot \mathbf{H}_G^{(L)}[k, :]$ where $\Phi(G)_y$ is the output logit for class y , and $n_G^k = \oplus_{j=1}^{|V|} \mathbf{H}_G^{(L)}[k, j]$ is the pooled value for neuron k in which \oplus is a global pooling operator. Now for η_E^k we aggregate the scaled node activations. Denoting $\boldsymbol{\eta}_V^k = (\eta_V^k, \eta_V^k, \dots, \eta_V^k) \in \mathbb{R}^{|V| \times |V|}$, we compute $\eta_E^k = (\boldsymbol{\eta}_V^k + (\boldsymbol{\eta}_V^k)^\top) \odot A$ where \odot is the Hadamard product and A is the adjacency matrix of G .

Now, for each neuron $k \in S_{\text{glob}}$ we choose $G = \arg \max_{G \in \mathcal{D}} \max_j \mathbf{H}_G[j, i]$ as the graph that has the highest presence of neuron k , and produce the concept activation map for (Φ, G, y, k) .

Contextualising global explanations using instances

The proposed method is able to extract concepts on a model-level. This is accompanied by a visualisation on an instance-level to aid the explanation. As there is growing interest in relating global explanations to local instances (Lundberg et al. 2020), we introduce metrics to assess how important a concept is in the model’s decision:

- *Absolute contribution* (ABS): the value of the neuron’s weight connection to the class output.
- *Entropy score* (ENT): we propose a mutual-information maximisation approach inspired by GNNEExplainer (Ying et al. 2019) to find the set of neurons S_{loc} which best explains a local prediction, whereby we optimise the objective

$$S_{\text{loc}} = \arg \min_S -\log P_\Phi(Y = y \mid \mathcal{X} = \sigma(\eta_N)), \quad (6)$$

in which P_Φ is the logit for class y , \mathcal{X} is the set of neurons utilised by the model, and η_N is a learned mask over the neurons. The full derivation is shown in Appendix B. The entropy importance of neuron k is mask value $\eta_N[k]$.

What Concepts Do GNNs Learn?

In this section we use the concept extraction approach to investigate the types of concepts that arise in GNNs, and demonstrate that they align with human intuition.

Datasets and models We perform experiments on four popular real-world benchmark datasets for binary graph classification: MUTAG (Debnath et al. 1991), Reddit-Binary (Yanardag and Vishwanathan 2015), PROTEINS (Borgwardt et al. 2005) and IMDB-Binary (Yanardag and Vishwanathan 2015). Following the existing trend in GNN explainability, we choose to train models which fit under the convolutional flavour of GNNs. For each graph classification task we train a model with three layers of the Graph Isomorphism Network (GIN) architecture (Xu et al. 2019). The exception is Reddit-Binary, where we use graph convolutional network (GCN) layers (Welling and Kipf 2016) as we observed training to be more stable given the large graph sizes. The exact descriptions of the models and their training parameters are shown in Appendix C. The same concept discovery procedure is run for all datasets, with a different set of base concepts for each. These were determined based

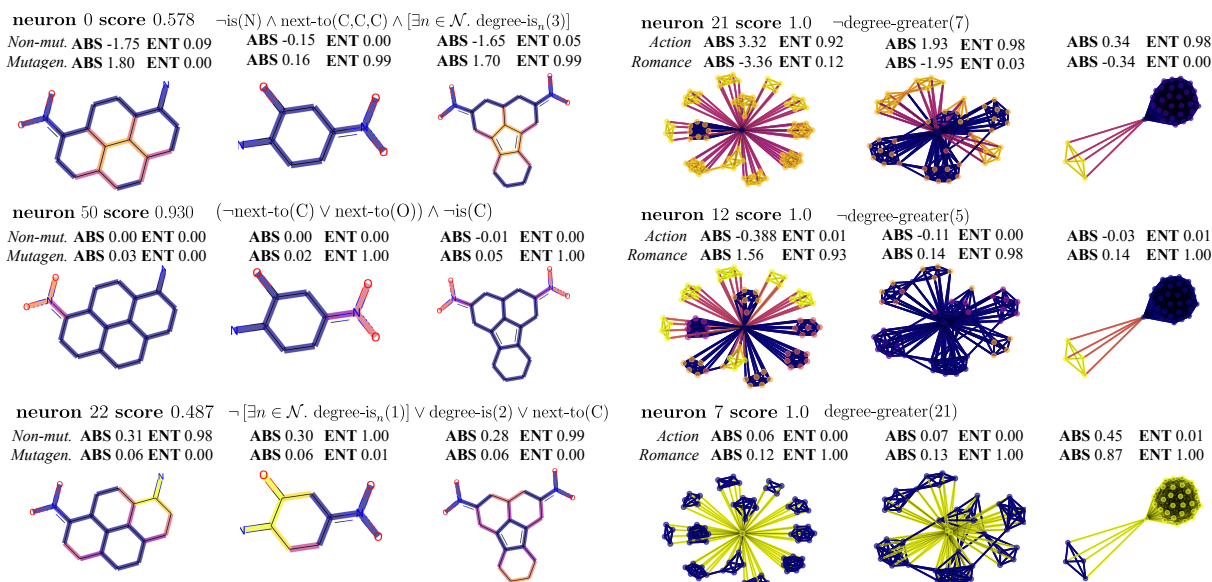


Figure 2: Concepts in MUTAG (left) and IMDB-B (right). Brighter colours indicate higher importance. Each row is thresholded using the threshold τ_k .

on the nature of the task. For example in MUTAG, the nodes can be one of seven types, and it is natural to include these as base concepts.

Concepts for Each Model

We apply our GNN concept discovery framework on the models trained on the benchmark datasets, using a set of base concepts for each task to align with the human-in-the-loop nature of explainability (Simkute et al. 2021). We choose *simple* concepts that non-domain experts are able to understand, such as *next-to*(X) (being next to a node with label X), and *is*(X) (having label X). The full set of base concepts for each task and their descriptions is shown in Appendix E. We find that interpretable concepts *are* detected by the neurons of a GNN, and they vary across the different tasks. Using concept activation maps we produce visual summaries of the most interpretable concepts. The concept contribution scores for each example graph using *absolute contribution* and *entropy score* are provided.

MUTAG We find that concepts in MUTAG correspond to known *functional groups*. As shown in Figure 2, we find NO_2 detector concepts and carbon ring concepts, which corroborate with the findings of Ying et al. (2019). The NO_2 concept is in fact a known mutagen which has been shown to be a strong indicator of mutagenicity (Kazius, McGuire, and Bursi 2005), and the presence of benzene carbon rings is an indicator of aromaticity which together with NO_2 is known to indicate mutagenicity. In the examples in Figure 2, **neuron 0** is a carbon ring detector and **neuron 50** is an NO_2 detector. We observe that carbon ring concepts typically have much higher absolute contribution to the final prediction. Unlike prior work, we additionally observe concepts which contribute to the prediction of non-mutagenicity. In particular we observe concepts that are highly concentrated

on parts of compounds that hang off carbon rings but are not NO_2 substructures - for instance, **neuron 22** appears to detect any part of the input which is not part of a carbon ring or an NO_2 substructure. Other concepts appear to trigger for only certain types of atoms hanging off carbon rings, e.g. **neuron 28** (not shown) detects atoms that are not fluorine nor chlorine, and have exactly one carbon neighbour.

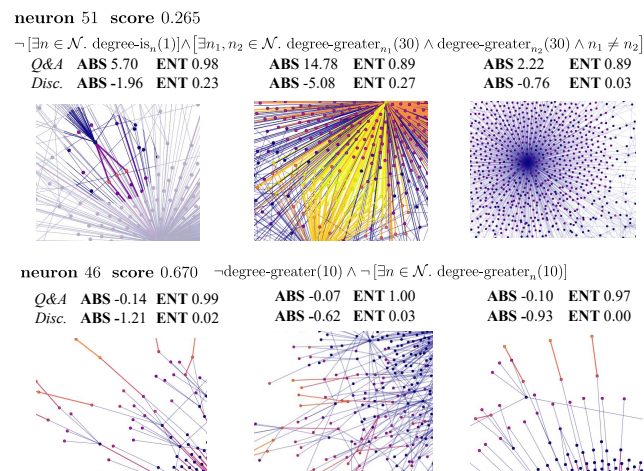


Figure 3: Reddit-Binary Concepts. Node colours represent activation values for the concept, with brighter colours indicating higher values. For graphs where the concept is hard to see by eye, we have highlighted the concept area by changing the opacity of the image.

Reddit-Binary Examples of concepts are shown in Figure 3. In Reddit-Binary, graphs typically have one or two nodes with high degree, corresponding to the highly voted comments that receive the most attention. We observe the emer-

gence of the ‘few-to-many’ subgraphs, resembling a small number of experts replying to several questions, as was mentioned by Ying et al. (2019), and is a strong indication that the graph is of the Q&A class. An example of such a concept detector is **neuron 51**. We also see neurons which appear to detect the nodes that are more than a single hop from a high-degree node, representing users at the end of a long comment chain, e.g. **neuron 46**. We observe that these concepts penalise the discussion class in terms of absolute contribution. One possible explanation is that discussion threads typically involve the same users commenting several times in the same chain, which would cause the chain to collapse in the graph (since nodes represent users, not comments). We additionally find neurons that appear to trigger when their degree is above a certain threshold, such as **neuron 20** (not shown). Intuitively, a user with lots of replies is a strong indication of Q&A, since the host of the Q&A will likely receive a large number of follow-up questions.

IMDB-Binary Graphs typically have one or two ‘popular’ nodes that have high degree, and have multiple smaller communities attached to them. As shown in Figure 2, we observe that almost all active neurons are sensitive to the *degree* of a node, with some neurons only activating if the degree exceeds (or is less than) a certain amount. We find that **neuron 21** is activated for the small communities that are attached to the popular node, and is accurately summarised as detecting if the degree is no greater than 7. It acts as a strong indicator of the genre being *action*. Interestingly, **neuron 12** behaves similarly - activating for nodes with degree no more than 5 - but is a strong indicator for the opposite genre *romance*. This suggests the GNN is learning to discriminate genres based on subtle differences in the size of the small communities. There exists high-degree concepts, e.g. **neuron 7** which are indicative of the romance class. Moreover, we are surprised by the high interpretability score of these neurons, with many having a perfect score of 1, i.e. the concepts learnt by the GNN can be summarised accurately using only degree information.

PROTEINS We provide additional findings with the PROTEINS dataset in Appendix F.

Discussion

Overall we observe that a range of concepts are being detected by GNN neurons. We make three observations.

1. We first observe that the extracted concepts are more interpretable than prior work for vision and language models. This is attributed to the fact that GNNs tend to be much shallower than CNNs and RNNs, and the ones used in the experiments have at most three layers. This is perhaps a hopeful sign from the perspective of GNN interpretability and the desirability for concepts to be extracted without resorting to manual labelling.
2. Secondly, we see that not all concepts are useful for the model. There exist neurons which detect concepts, but have low contribution to predictions. In addition, there exist neurons that the model believes to be useful for

both classes and therefore do not act as useful discriminators for the model. For example, in IMDB-Binary, there are neurons that detect the presence of a node with large degree which have positive weights for both *romance* and *action* genres. This suggests that by manually removing or suppressing certain concepts we may be able to reduce the number of parameters whilst preserving performance, which serves as a form of *model pruning*. We leave this idea open for future investigation.

3. Moreover, there are often neurons that detect similar concepts, and appear to be redundant in nature. For instance, the MUTAG model contains at least 3 neurons that act more or less as NO₂ detectors, but their associated concepts are marginally different. This can be attributed to the nature of SGD-based algorithms used to train GNNs - there is no explicit constraint that enforces each neuron to behave differently. We note that despite there being many neurons detecting similar concepts, their activation magnitudes can still vary greatly. This suggests the method can be potentially used to identify neurons that are redundant in the information they learn.

Are Interpretable Concepts More Important?

A natural question to ask is whether the concepts identified are beneficial to the model making correct predictions. We define a neuron’s *interpretability* as the score of its best matching concept. Existing methods of determining neuron importance do not necessarily work for GNNs as they do for CNNs. For example, we found that *erasure* (Li, Monroe, and Jurafsky 2016) is a poor metric for determining neuron importance due to the high presence of *redundant neurons* learnt by GNNs. We propose two alternative metrics.

Neuron importance. For a model with r class outputs, we measure the importance of neuron k in making a single prediction as the variance of individual logit contributions. Recall that for input $G = (E, V)$, $n_G^k = \bigoplus_{j=1}^{|V|} \mathbf{H}_G^{(L)}[k, j]$ is the pooled value of neuron k . Aiming to measure the neuron’s ability to discriminate between classes, rather than its absolute contribution, we consider the variance of individual logit contributions, i.e. $\text{Var}(w_k^1 n_G^k, \dots, w_k^r n_G^k)$ where w_k^j is the weight connecting n_G^k to the j -th logit. The *importance* score is this value averaged over all graphs in the dataset:

$$\text{importance}(k) = \frac{1}{|\mathcal{D}|} \sum_{G \in \mathcal{D}} \text{Var}(w_k^1 n_G^k, \dots, w_k^r n_G^k). \quad (7)$$

Neuron correctness. Following (Mu and Andreas 2020) we also measure how much each neuron contributes to the model being *accurate*. They chose to measure the frequency of correct predictions when each neuron is activated³. Let y_G be the predicted logits for input G and \hat{y}_G be the ground truth. Treating n_G^k , y_G , and \hat{y}_G as random variables sampled across $G \in \mathcal{D}$, we choose to compute a correctness score for each neuron k defined as

$$\text{correctness}(k) = |\text{Corr}(n_G^k, -\mathcal{L}(y_G, \hat{y}_G))|, \quad (8)$$

³While this yielded good results for CNNs, it does not transfer well to GNNs due to the redundant nature of neurons

where we measure the Pearson product-moment correlation between the neuron’s magnitude with cross-entropy loss.

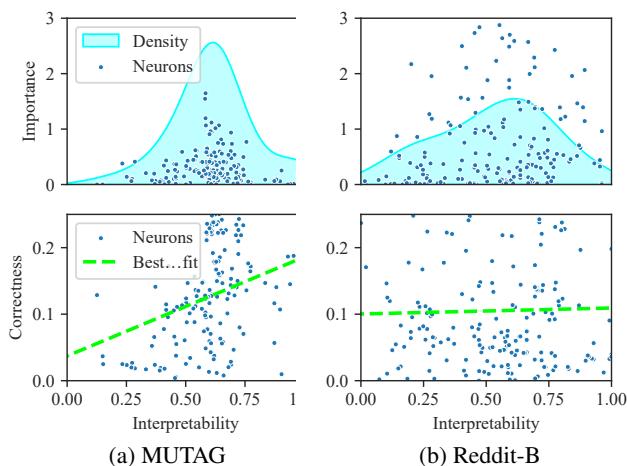


Figure 4: Investigating the interpretability of neurons and their importance. Kernel density estimates (Chen 2017) are shown in the top row to indicate the relative spread of neurons across different interpretability scores. Dashed lines are lines-of-best-fit to indicate trend.

We investigate the relationship between interpretability and the aforementioned metrics, by training ten models for each dataset across different train-test splits⁴. We show the results of the analysis in Figure 4. First, for the MUTAG and Reddit-Binary tasks, the distribution of neuron interpretability scores appears to resemble a bell curve, where more neurons have scores closer to 0.5 than either extreme. In terms of the relationship with neuron importance, we observe that the most important neurons are situated around the centre of the interpretability spectrum. This suggests that the neurons that are most important are also more difficult to summarise using a set of interpretable rules. Nevertheless, we do not observe a negative correlation between neuron importance and interpretability, and the most important neurons typically have around 0.5 interpretability, which we empirically observe to be sufficient in terms of explanation quality. In terms of contribution to model accuracy, we observe that neuron interpretability is positively correlated with having accurate predictions for MUTAG, and having slight positive correlations for Reddit-Binary.

How Do Concepts Vary Across Models?

Works on interpreting deep models with concepts typically focus on well-trained models that generalise well. They also focus on the last few hidden layers since high-level concepts are more likely to emerge there. This motivates us to ask two follow-up questions:

1. Do models with only a few epochs of training show different concepts than models trained for longer?
2. Do different concepts emerge at different depths?

⁴IMDB-Binary is excluded since almost all neurons are highly interpretable

We perform additional experiments on the MUTAG and Reddit-Binary datasets in an attempt to answer these questions and provide some preliminary steps for future investigation.

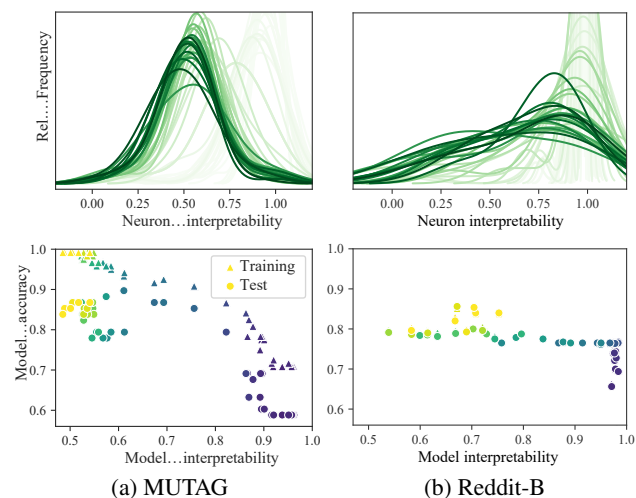


Figure 5: Training duration and model interpretability. The top row shows the distribution of neuron interpretability scores for models trained for varying epochs; darker curves correspond to more epochs. The lower row shows the relationship between the interpretability of a model and its accuracy; brighter points correspond to more epochs.

Longer training reduces interpretability. We train models using the same architectures as previously specified, but for varying numbers of epochs ranging from 1 to 400. For each model, we extract concepts from the final layer. Defining *model interpretability* as the mean interpretability of its neurons, we observe the interpretability distributions for these models, as well as the relationship between model interpretability and accuracy. This is shown in Figure 5. Training a model for longer periods causes concepts to be less interpretable, in general. Note that in some cases, there is a small difference in test set performance, despite a large difference in concept interpretability. For example, both the MUTAG models trained for 100 epochs and 400 epochs have the same test accuracy, but one has significantly more interpretable concepts. This phenomenon is more extreme on Reddit-Binary; we observe that for the first 50 epochs the model performance improves without much sacrifice in interpretability; however, later the interpretability drops significantly without much improvement in model performance. This suggests that there may be some general trade-off between interpretability and model accuracy, even when the same model architecture is used. Therefore, interpretability could be a factor to be considered during model training. A related theoretical issue is the *information bottleneck* in deep learning (Shwartz-Ziv and Tishby 2017), which makes the observation that there is a trade-off between accuracy and the information flow in the latter layers of a model.

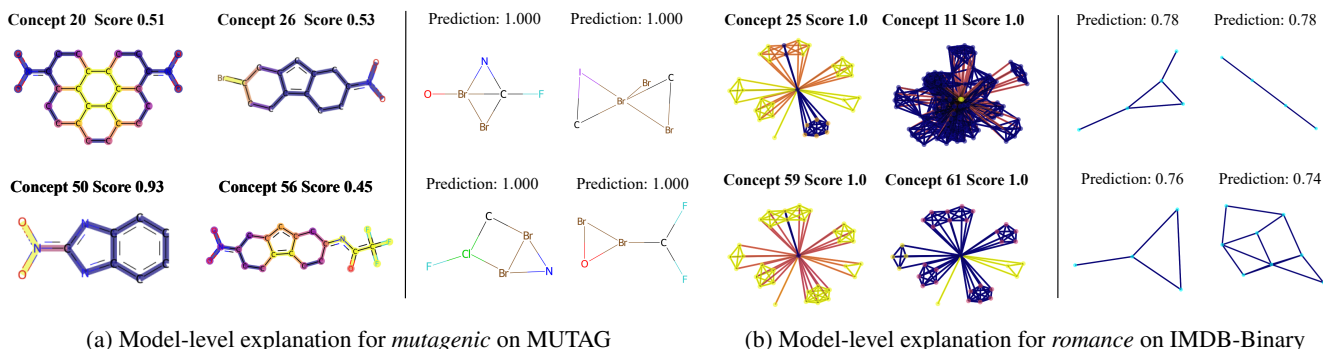


Figure 6: Side-by-side comparison of class-specific model-level explanations against XGNN on MUTAG and IMDB-Binary. In each subfigure, our outputs are shown in the left; XGNN outputs are on the right along with the prediction values.

Concepts across layers. We train additional models on the MUTAG and Reddit-Binary datasets, but using 10 layers of message passing. We probe the concepts after each layer and find that the mean neuron interpretability decreases further into the model. Refer to Appendix D for plots. The accuracy of these deeper GNNs is lower than the shallower versions, which is not surprising and likely due to *over-smoothing* (Li, Han, and Wu 2018). However, we also notice that the interpretability of the final layer is typically lower than those of the shallow models, suggesting that there is also a potential link between over-smoothing and neuron interpretability.

Model-Level Explanations

In this section, we provide concept-based explanations using our model-level approach. We compare the explanations against the current state-of-the-art approach for model-level explanations for GNNs, XGNN (Yuan et al. 2020; Dai et al. 2022), which produces class-specific global explanations via *graph generation*. We compare against XGNN in terms of the model-level explanations produced by both methods, showcasing the concepts of our method against the generated graphs of XGNN. Comparisons on MUTAG and IMDB-Binary are shown in Figure 6; additional comparisons on Reddit-Binary are provided in Appendix G.

Comparison We want to emphasise that our approach offers desirable properties which XGNN is not capable of. In particular, it can be observed that:

- Our approach is capable of separating the explanation into *individual concepts*, which has been shown to align with human intuition (Koh et al. 2020). XGNN is not concept-based and has no guarantee that each generated graph corresponds to a separate unit of decision making.
- XGNN’s generated graphs leave the human observer to draw a conclusion. Our approach, however, can provide logical descriptions of each concept together with an interpretability score. We argue that this reduces the potential for human bias compared to providing only a generated graph, which can have very different interpretations depending on the observer.

- XGNN often generates graphs that do not resemble members of the actual data distribution, making it difficult to contextualise the explanation in the problem setting. For example, we do not see the appearance of carbon-ring concepts for the MUTAG class *mutagenic*. Similarly, we do not see the *few-to-many* concept for Reddit-Binary. For IMDB-Binary, from the XGNN outputs we cannot infer that the class prediction is related to node degree.
- We are able to contextualise our model-level explanations using local instances using the importance metrics we proposed, which is not possible with XGNN.
- Our approach avoids the ‘explaining a black-box via another black-box’ issue (Kovalerchuk, Ahmad, and Tere-desai 2021).

Usage of XGNN We use the implementation of XGNN provided in DIG (Liu et al. 2021). We remark that XGNN is able to generate graphs maximising the prediction score on MUTAG, as shown in (Yuan et al. 2020). However, we find that XGNN fails to generate a graph with score above 0.8 on IMDB *romance* class, even after thorough tuning of parameters. For the *action* class, we find that XGNN fails to generate a graph beyond a single node.

Conclusions

We perform a new explorative study of the presence of concepts in graph neural networks and provide insights which have potential to improve our understanding of GNNs and help practitioners gain a more transparent view of their models. To the best of our knowledge this is the first work that tackles GNN interpretability via neuron-level analysis of GNN representations. Furthermore, we use concepts to provide model-level explanations, which show promising advantages over the current state-of-the-art. This is only an initial step in unveiling the neuron-level concepts that exist in GNNs, and much work can be done in the directions we have outlined in the paper. A promising future direction is to further explore the relationship between neuron interpretability and theoretical issues such as the information bottleneck, over-smoothing and over-squashing.

References

- Amann, J.; Blasimme, A.; Vayena, E.; Frey, D.; and Madai, V. I. 2020. Explainability for Artificial Intelligence in Healthcare: A Multidisciplinary Perspective. *BMC Medical Informatics and Decision Making*, 20(1): 1–9.
- Barbiero, P.; Ciravegna, G.; Giannini, F.; Liò, P.; Gori, M.; and Melacci, S. 2022. Entropy-based logic explanations of neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6046–6054.
- Bau, D.; Zhou, B.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Network Dissection: Quantifying Interpretability of Deep Visual Representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6541–6549.
- Borgwardt, K. M.; Ong, C. S.; Schönauer, S.; Vishwanathan, S. V. N.; Smola, A. J.; and Kriegel, H.-P. 2005. Protein Function Prediction via Graph Kernels. *Bioinformatics*.
- Bussmann, N.; Giudici, P.; Marinelli, D.; and Papenbrock, J. 2021. Explainable Machine Learning in Credit Risk Management. *Computational Economics*, 57(1): 203–216.
- Chen, Y.-C. 2017. A Tutorial on Kernel Density Estimation and Recent Advances. *Biostatistics & Epidemiology*, 1(1): 161–187.
- Dai, E.; Zhao, T.; Zhu, H.; Xu, J.; Guo, Z.; Liu, H.; Tang, J.; and Wang, S. 2022. A Comprehensive Survey on Trustworthy Graph Neural Networks: Privacy, Robustness, Fairness, and Explainability. *arXiv preprint arXiv:2204.08570*.
- Dalvi, F.; Durrani, N.; Sajjad, H.; Belinkov, Y.; Bau, A.; and Glass, J. 2019. What Is One Grain of Sand in the Desert? Analyzing Individual Neurons in Deep NLP Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6309–6317.
- Debnath, A. K.; Lopez de Compadre, R. L.; Debnath, G.; Shusterman, A. J.; and Hansch, C. 1991. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of Medicinal Chemistry*, 34(2): 786–797.
- Deeks, A. 2019. The Judicial Demand for Explainable Artificial Intelligence. *Columbia Law Review*, 119(7): 1829–1850.
- Georgiev, D.; Barbiero, P.; Kazhdan, D.; Veličković, P.; and Liò, P. 2022. Algorithmic Concept-based Explainable Reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6685–6693.
- Ghorbani, A.; Wexler, J.; Zou, J. Y.; and Kim, B. 2019. Towards Automatic Concept-Based Explanations. *Advances in Neural Information Processing Systems*, 32.
- Himmelhuber, A.; Zillner, S.; Grimm, S.; Ringsquandl, M.; Joblin, M.; and Runkler, T. A. 2021. A New Concept for Explaining Graph Neural Networks. In *NeSy*, 1–5.
- Kazius, J.; McGuire, R.; and Bursi, R. 2005. Derivation and Validation of Toxicophores for Mutagenicity Prediction. *Journal of Medicinal Chemistry*, 48(1): 312–320.
- Kim, B.; Wattenberg, M.; Gilmer, J.; Cai, C.; Wexler, J.; Viegas, F.; et al. 2018. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International Conference on Machine Learning*, 2668–2677. PMLR.
- Koh, P. W.; Nguyen, T.; Tang, Y. S.; Mussmann, S.; Pierson, E.; Kim, B.; and Liang, P. 2020. Concept Bottleneck Models. In *International Conference on Machine Learning*, 5338–5348. PMLR.
- Kovalerchuk, B.; Ahmad, M. A.; and Teredesai, A. 2021. Survey of Explainable Machine Learning with Visual and Granular Methods Beyond Quasi-Explanations. *Interpretable Artificial Intelligence: A Perspective of Granular Computing*, 217–267.
- Li, J.; Monroe, W.; and Jurafsky, D. 2016. Understanding Neural Networks Through Representation Erasure. *arXiv preprint arXiv:1612.08220*.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper Insights into Graph Convolutional Networks for Semi-Supervised Learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Liu, M.; Luo, Y.; Wang, L.; Xie, Y.; Yuan, H.; Gui, S.; Yu, H.; Xu, Z.; Zhang, J.; Liu, Y.; et al. 2021. DIG: A Turnkey Library for Diving into Graph Deep Learning Research. *Journal of Machine Learning Research*, 22(240): 1–9.
- Lundberg, S. M.; Erion, G.; Chen, H.; DeGrave, A.; Prutkin, J. M.; Nair, B.; Katz, R.; Himmelfarb, J.; Bansal, N.; and Lee, S.-I. 2020. From Local Explanations to Global Understanding with Explainable AI for Trees. *Nature Machine Intelligence*, 2(1): 56–67.
- Luo, D.; Cheng, W.; Xu, D.; Yu, W.; Zong, B.; Chen, H.; and Zhang, X. 2020. Parameterized Explainer for Graph Neural Network. *Advances in Neural Information Processing Systems*, 33.
- Magister, L. C.; Kazhdan, D.; Singh, V.; and Liò, P. 2021. GCEExplainer: Human-in-the-Loop Concept-based Explanations for Graph Neural Networks. *arXiv preprint arXiv:2107.11889*.
- Mu, J.; and Andreas, J. 2020. Compositional Explanations of Neurons. *Advances in Neural Information Processing Systems*, 33: 17153–17163.
- Samek, W.; Montavon, G.; Vedaldi, A.; Hansen, L. K.; and Müller, K.-R. 2019. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700. Springer Nature.
- Sanchez-Gonzalez, A.; Godwin, J.; Pfaff, T.; Ying, R.; Leskovec, J.; and Battaglia, P. 2020. Learning to Simulate Complex Physics with Graph Networks. In *International Conference on Machine Learning*, 8459–8468. PMLR.
- Selvaraju, R. R.; Das, A.; Vedantam, R.; Cogswell, M.; Parikh, D.; and Batra, D. 2019. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal of Computer Vision*, 128: 336–359.
- Shwartz-Ziv, R.; and Tishby, N. 2017. Opening the Black Box of Deep Neural Networks via Information. *arXiv preprint arXiv:1703.00810*.

- Simkute, A.; Luger, E.; Jones, B.; Evans, M.; and Jones, R. 2021. Explainability for Experts: A Design Framework for Making Algorithms Supporting Expert Decisions More Explainable. *Journal of Responsible Technology*, 7-8: 100017.
- Veličković, P. 2022. Message Passing All the Way Up. *arXiv preprint arXiv:2202.11097*.
- Vu, M.; and Thai, M. T. 2020. PGM-Explainer: Probabilistic Graphical Model Explanations for Graph Neural Networks. *Advances in Neural Information Processing Systems*, 33: 12225–12235.
- Welling, M.; and Kipf, T. N. 2016. Semi-Supervised Classification with Graph Convolutional Networks. In *J. International Conference on Learning Representations (ICLR 2017)*.
- Wieder, O.; Kohlbacher, S.; Kuenemann, M.; Garon, A.; Ducrot, P.; Seidel, T.; and Langer, T. 2020. A Compact Review of Molecular Property Prediction with Graph Neural Networks. *Drug Discovery Today: Technologies*, 37: 1–12.
- Wu, L.; Chen, Y.; Shen, K.; Guo, X.; Gao, H.; Li, S.; Pei, J.; and Long, B. 2021. Graph Neural Networks for Natural Language Processing: A Survey. *arXiv preprint arXiv:2106.06090*.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? *ArXiv*, abs/1810.00826.
- Yanardag, P.; and Vishwanathan, S. 2015. Deep Graph Kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1365–1374.
- Ying, Z.; Bourgeois, D.; You, J.; Zitnik, M.; and Leskovec, J. 2019. GNNExplainer: Generating Explanations for Graph Neural Networks. *Advances in Neural Information Processing Systems*, 32.
- Yuan, H.; Tang, J.; Hu, X.; and Ji, S. 2020. XGNN: Towards Model-Level Explanations of Graph Neural Networks. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning Deep Features for Discriminative Localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2921–2929.