

Table 3: The best hyperparameter configurations achieved for each dataset using the GIN model.

Dataset	NUM_LAYERS_PRE	NUM_LAYERS	DIM_INNER	DROPOUT	BASE_LR
BZR	1	2	32	0.0	$1e^{-3}$
COX2	2	2	128	0.0	$1e^{-2}$
DD	2	2	128	0.5	$1e^{-3}$
ENZYMES	2	1	128	0.5	$1e^{-3}$
MUTAG	2	1	128	0.0	$1e^{-2}$
NCI1	1	2	128	0.0	$1e^{-3}$
NCI109	2	2	128	0.0	$1e^{-3}$
PROTEINS	1	2	128	0.5	$1e^{-3}$
PTC	1	2	128	0.0	$1e^{-3}$

## A TRAINING DETAILS

In this section, we provide detailed information regarding the experimental setup used to obtain the results. We begin by describing the hyperparameters that were cross-validated for the GNN graph-level classifier on GIN. The optimal configuration obtained for each dataset was then used to train the rest of the models. We also present the additional hyperparameters that were cross-validated for the other models apart from GIN. All the fixed hyperparameters can be found in the GitHub repository at <https://github.com/XXXX/XXXX>. For all the experiments, we trained the models using a seed of 42 and report the mean and standard deviation over 5 different dataset splits. We run the models based on reinforcement learning, i.e., SUGAR (<https://github.com/RingBDStack/SUGAR>) and GCIP, for 1000 epochs and the rest of the models for 500 epochs. The experiments were conducted on a single CPU with 10GB RAM.

**GIN Hyperparameters** The following hyperparameters were cross-validated with the corresponding values:

- NUM\_LAYERS\_PRE: Number of layers of the MLP used prior to the GNN. Cross-validated values: [1, 2].
- NUM\_LAYERS: Number of layers of the GNN. Cross-validated values: [1, 2].
- DIM\_INNER: Number of neurons in each (graph) neural network layer. Cross-validated values: [32, 128].
- DROPOUT: Probability of applying dropout. Cross-validated values: [0.0, 0.5].
- BASE\_LR: Base learning rate used for training. Cross-validated values:  $\{1e^{-2}, 1e^{-3}, 1e^{-4}\}$ .

Table 3 presents the best configuration obtained for each dataset, which was subsequently used to train the remaining models.

**TopK Hyperparameters** The following hyperparameter was cross-validated with the corresponding values:

- RATIO: Percentage of nodes to retain. Cross-validated values: {0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9}.

**RNDN Hyperparameters** The following hyperparameter was cross-validated with the corresponding values:

- RATIO: Percentage of nodes to retain. Cross-validated values: {0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9}.

**RNDE Hyperparameters** The following hyperparameter was cross-validated with the corresponding values:

Table 4: The best hyperparameter configurations achieved for each dataset using the GCIP model for node (left) and edge (right) removal.  $d$  represents the maximum desired node/edge ratio.

Dataset	GCIP <sub>N</sub>				GCIP <sub>E</sub>			
	$\lambda$	$d$	SPLITS	BASE_LR_2	$\lambda$	$d$	SPLITS	BASE_LR_2
BZR	0.50	0.30	[0.5, 0.4, 0.1]	1e-4	0.50	0.30	[0.6, 0.3, 0.1]	1e-4
COX2	0.50	0.30	[0.6, 0.3, 0.1]	1e-4	0.00	0.25	[0.5, 0.4, 0.1]	1e-4
DD	0.00	0.95	[0.6, 0.3, 0.1]	1e-4	0.50	0.30	[0.6, 0.3, 0.1]	1e-4
ENZYMES	1.00	0.30	[0.5, 0.4, 0.1]	1e-4	0.50	0.30	[0.5, 0.4, 0.1]	1e-3
MUTAG	0.50	0.75	[0.5, 0.4, 0.1]	1e-3	1.00	0.30	[0.6, 0.3, 0.1]	5e-3
NCI1	0.50	0.30	[0.6, 0.3, 0.1]	1e-4	0.00	0.50	[0.5, 0.4, 0.1]	1e-4
NCI109	0.00	0.75	[0.6, 0.3, 0.1]	1e-4	0.25	0.30	[0.6, 0.3, 0.1]	1e-4
PROTEINS	1.00	0.30	[0.5, 0.4, 0.1]	5e-3	0.75	0.30	[0.6, 0.3, 0.1]	5e-3
PTC	0.00	0.50	[0.5, 0.4, 0.1]	1e-4	0.00	0.95	[0.6, 0.3, 0.1]	1e-3

- **RATIO:** Percentage of edges to retain. Cross-validated values: {0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9}.

**DiffPool Hyperparameters** The following hyperparameters were cross-validated with the corresponding values:

- **ADDITIONAL\_LOSSES:** Whether or not to include the additional loss proposed in Ying et al. (2018). Cross-validated values: {True, False}.
- **BASE\_LR:** Base learning rate used for training. Cross-validated values: { $1e^{-2}$ ,  $1e^{-3}$ ,  $1e^{-4}$ }.

**SUGAR Hyperparameters** The following hyperparameters were cross validated with the corresponding values:

- **NEGATIVE\_SAMPLING\_RATIO:** This hyperparameter was cross validated for the self-supervised MI Module and it controlled the amount of negative subgraphs being sampled for contrastive learning. Cross-validated values {1, 2, 3, 4, 5}
- **BETA:** This hyperparameter acted as coefficient for the self supervised module in SUGAR. Cross-validated values {0.0, 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4, 1.6, 1.8, 2.0}

**GCIP Hyperparameters** The following hyperparameters were cross-validated with the corresponding values. The same values were cross-validated for both GCIP<sub>N</sub> and GCIP<sub>E</sub>. Table 4 presents the best configuration obtained for each dataset for both GCIP<sub>N</sub> and GCIP<sub>E</sub>.

- **BASE\_LR\_2:** Base learning rate used for the outer optimization, which corresponds to the reinforcement learning component. Cross-validated values: { $1e^{-3}$ ,  $1e^{-4}$ }.
- **SPLITS:** The split percentages used for training, validation, and test sets. The test size was kept fixed, while the training and validation sizes were varied. These splits were used for training the base graph classifier and the reinforcement learning component, respectively. Cross-validated values: {[0.5, 0.4, 0.1], [0.6, 0.3, 0.1]}.

#### A.1 EVOLUTION OVER EPOCHS

We present (in Figure 5) the evolution of several metrics during the training of the nine datasets discussed in Section 5. Specifically, we analyze the accuracies, PPO rewards, node/edge ratio.

Upon observing the trends across all metrics, we find a consistent and stable evolution throughout the epochs, providing compelling evidence for the effectiveness of the bi-level optimization approach. However, it is important to note that three datasets (NCI1, NCI109, and PTC) do not achieve convergence by epoch 1000. This suggests that extending the training duration may further improve the performance of our proposed method GCIP. Nevertheless, to ensure a fair comparison, we did not pursue this approach.

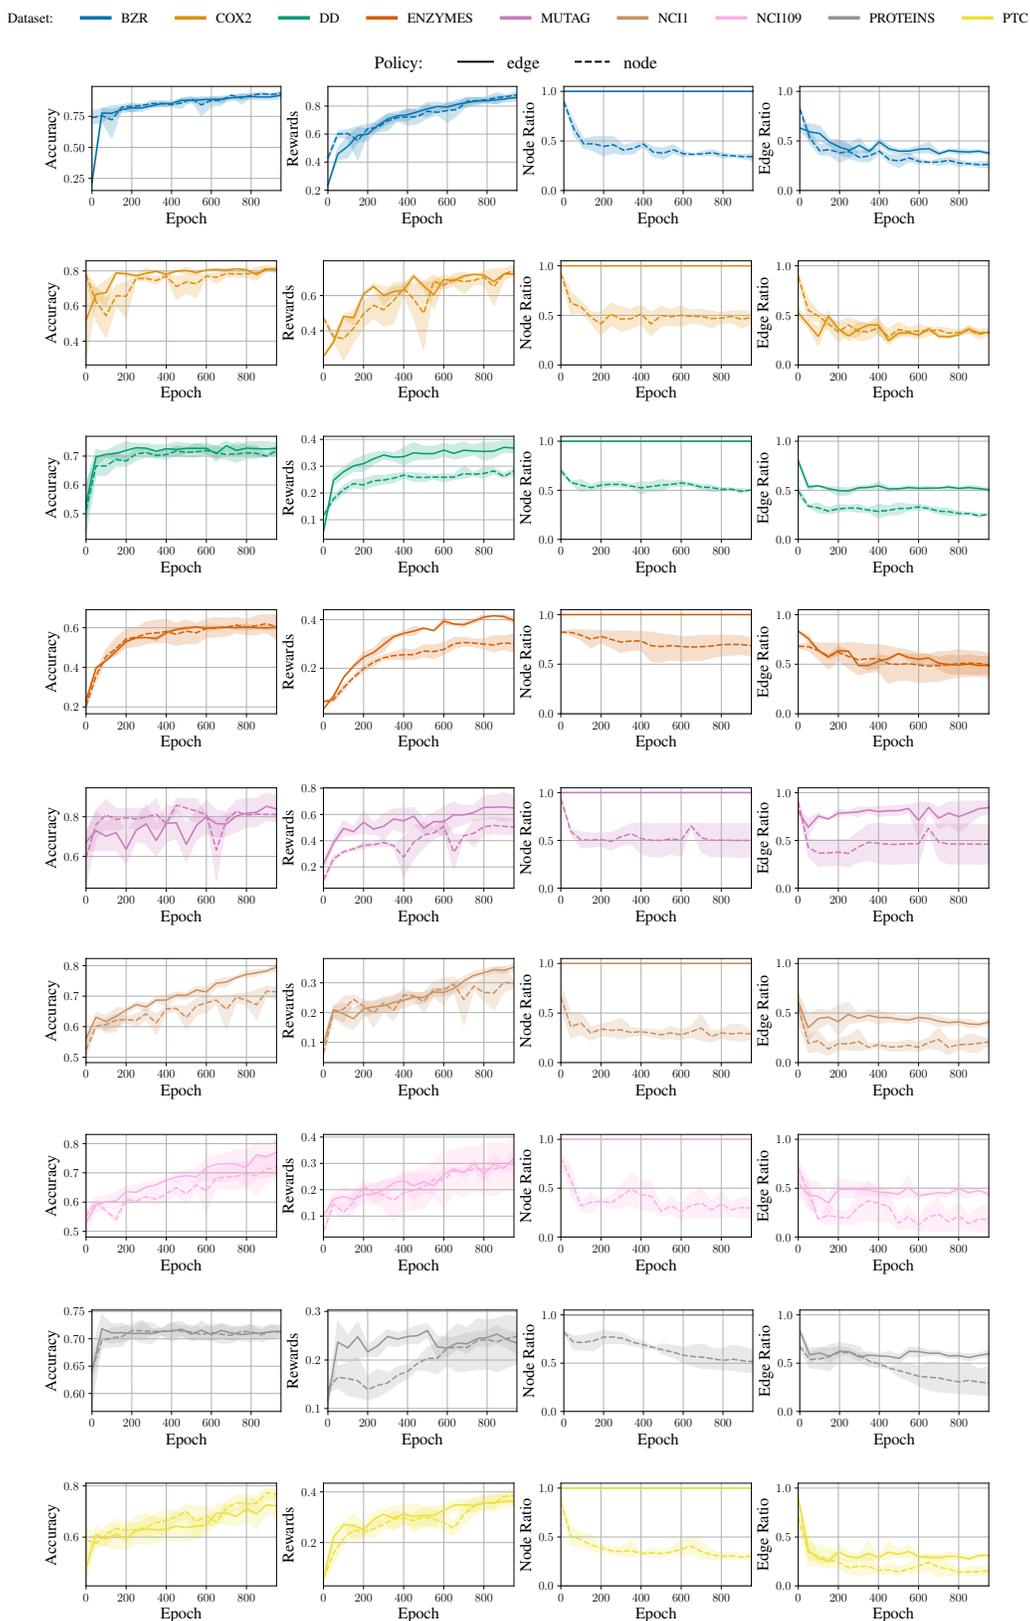


Figure 5: Evolution of different metrics during training for the nine datasets under study.

## B GCIP TRAINING PROCEDURE

In this section we include a description of the training procedure of GCIP, summarized in Algorithm 1.

---

**Algorithm 1** Training of GCIP.

---

```

1: Input: Training  $\mathcal{D}^{tr}$  and validation  $\mathcal{D}^{val}$  sets; Initial graph classifier parameters  $\theta$  and policy
   parameters  $\phi$ ; Empty buffer  $\mathcal{B}$ ; Learning rates  $\alpha$  and  $\beta$ ; Number of PPO updates  $K$ 
2: while not convergence do
3:   for  $y, \mathcal{G} \in \mathcal{D}^{tr}$  do                                     ▷ Get a (batch of) sample(s) from the training set
4:      $\mathbf{a} \sim \pi_\phi(\mathbf{a}|\mathcal{G})$                                        ▷ Sample the nodes/edges to be removed using the policy
5:      $\mathcal{G}_s \leftarrow \mathcal{G}$                                            ▷ Get the subgraph
6:      $\hat{\mathbf{y}}_s = f_\theta(\mathcal{G}_s)$                                        ▷ Get the prediction
7:      $\theta \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}_{\text{perf}}(\theta, \phi)$        ▷ Update the parameters of the graph classifier
8:   end for
9:   for  $y, \mathcal{G} \in \mathcal{D}^{val}$  do                                       ▷ Get a (batch of) sample(s) from the validation set
10:     $\mathcal{G}_s \leftarrow \pi_\phi$                                            ▷ Use the  $\pi$  to sample a sparse graph
11:     $\hat{\mathbf{y}}_s = f(\mathcal{G}_s)$                                              ▷ Get the prediction of the sparse graph
12:    Add tuple  $\{y, \mathcal{G}, \hat{\mathbf{y}}, \mathcal{G}_s\}$  to the buffer  $\mathcal{B}$ 
13:  end for
14:  Compute rewards and prepare mini-batches from buffer  $\mathcal{B}$  for PPO updates
15:  for  $i = 1, 2, \dots, K$  do                                       ▷ Update PPO  $K$  times
16:     $\{y, \mathcal{G}, \hat{\mathbf{y}}, \mathcal{G}_s, r\} \sim \mathcal{B}$                                ▷ Sample a batch of tuples from the buffer
17:    Compute advantage estimates  $\hat{A}$ 
18:     $\phi \leftarrow \phi - \beta \nabla_\phi \mathcal{L}_{\text{spa}}(\theta, \phi)$            ▷ Update policy by maximizing Equation (1)
19:  end for
20:  Check for convergence, update convergence flag
21: end while
22: Output: Optimal policy parameters  $\phi$  and graph classifier parameters  $\theta$ 

```

---

## C COMPLETE RESULTS ON ABLATION STUDY

In this section, we provide a complete ablation study on  $\lambda$  and desired node/edge ratio  $d$  for all the datasets namely BZR, COX2, DD, ENZYMES, MUTAG, NCI1, NCI109, PROTEINS and PTC.

### C.1 ABLATION STUDY ON $\lambda$

Here we analyze the results of how the parameter  $\lambda$  affects the sparsity and performance of GCIP<sub>E</sub> and GCIP<sub>N</sub>, on all datasets. Results are shown in Figure 6 for GCIP<sub>N</sub> and Figure 7 for GCIP<sub>E</sub>.

We can observe that the  $\lambda$  is able to control sparsity in nodes and edges without dropping the accuracy. In general, lower values of  $\lambda$  result in sparser graphs. The effect is really clear for most datasets like NCI1, with the node ratio in Figure 6 ranging from approximately 0.5 to 0.15, and the edge ratio in Figure 7 ranging from approximately 0.8 to 0.4. However, for other datasets like DD, the differences are much smaller, with the node ratio in Figure 6 ranging from approximately 0.48 to 0.4, and the edge ratio in Figure 7 ranging from approximately 0.58 to 0.5.

### C.2 ABLATION STUDY ON MAXIMUM DESIRED RATIO $d$

Here we analyze the results of how the parameter  $d$  affects the sparsity and performance of GCIP<sub>E</sub> and GCIP<sub>N</sub>, on all datasets. Results are shown in Figure 8 for GCIP<sub>N</sub> and Figure 9 for GCIP<sub>E</sub>.

Similar to the results shown in Figure 6 and Figure 7, we can observe that the parameter  $d$  is able to control sparsity in nodes and edges without dropping the accuracy, extensively. Some datasets, such as DD, exhibit more subtle differences. However, in general, lower values of  $d$  clearly lead to sparser graphs. This is evident in datasets like NCI1, NCI109, and PTC, where the graph becomes significantly sparser as  $d$  decreases.

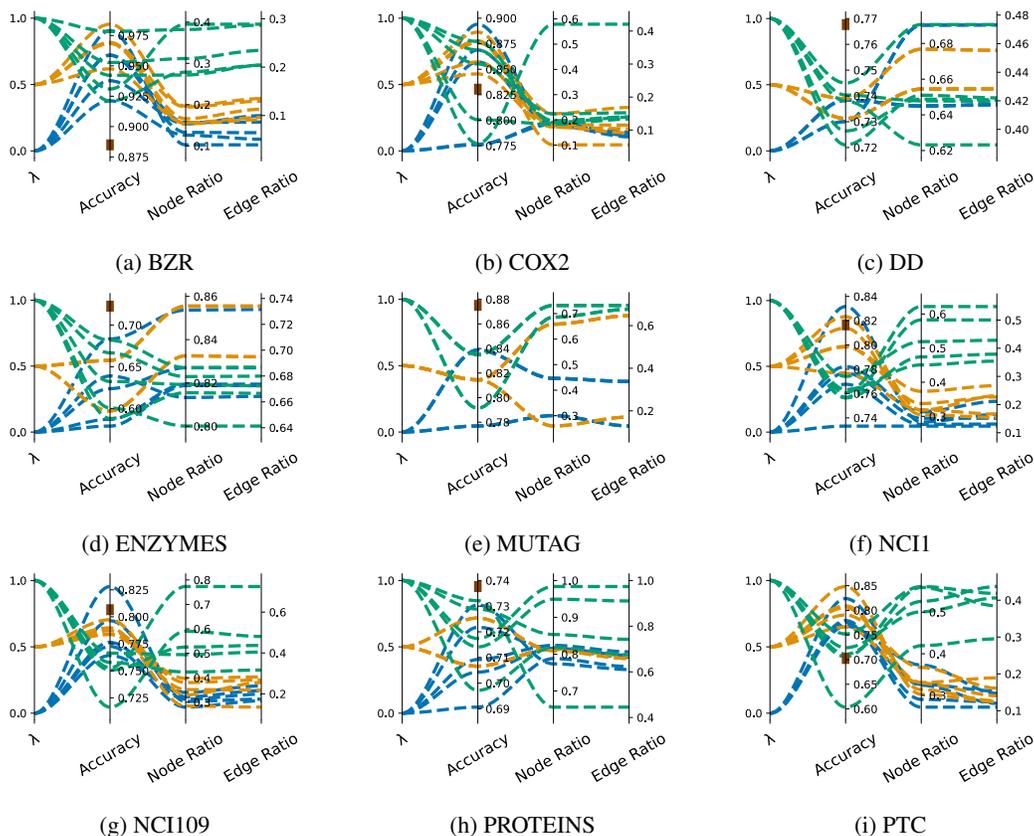


Figure 6: Ablation study on  $\lambda$  for  $GCIP_N$ , controlling the importance given to performance or sparsity. The brown solid rectangle on the accuracy axis represents the baseline GIN (baseline) accuracy.

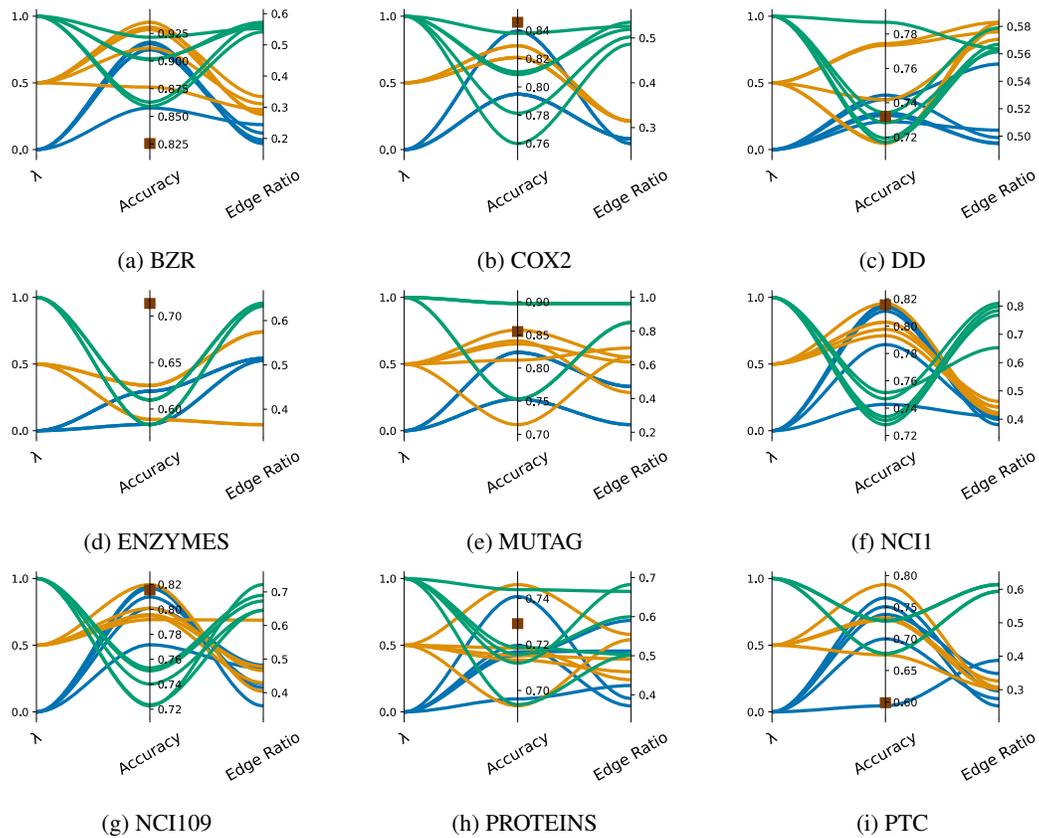


Figure 7: Ablation study on  $\lambda$  for  $GCIP_E$ , controlling the importance given to performance or sparsity. The brown solid rectangle on the accuracy axis represents the baseline GIN (baseline) accuracy.

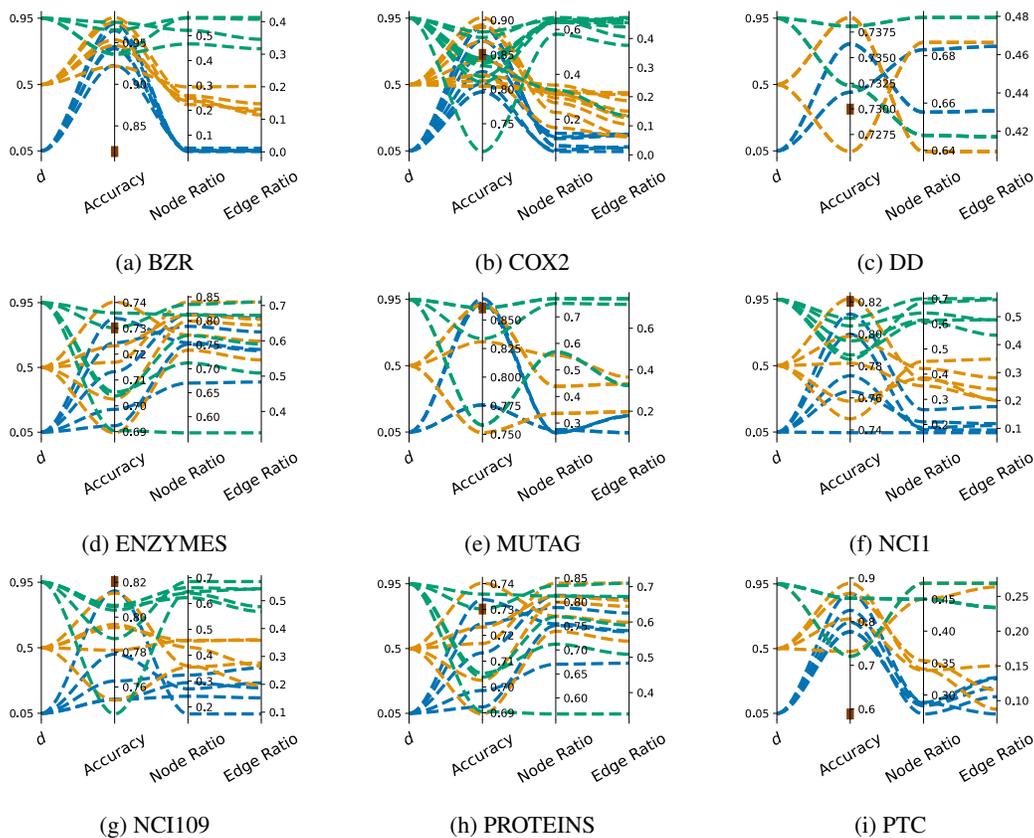


Figure 8: Ablation study on  $d$  for  $GCIP_N$ , controlling the importance given to performance or sparsity. The brown solid rectangle on the accuracy axis represents the baseline GIN (baseline) accuracy.

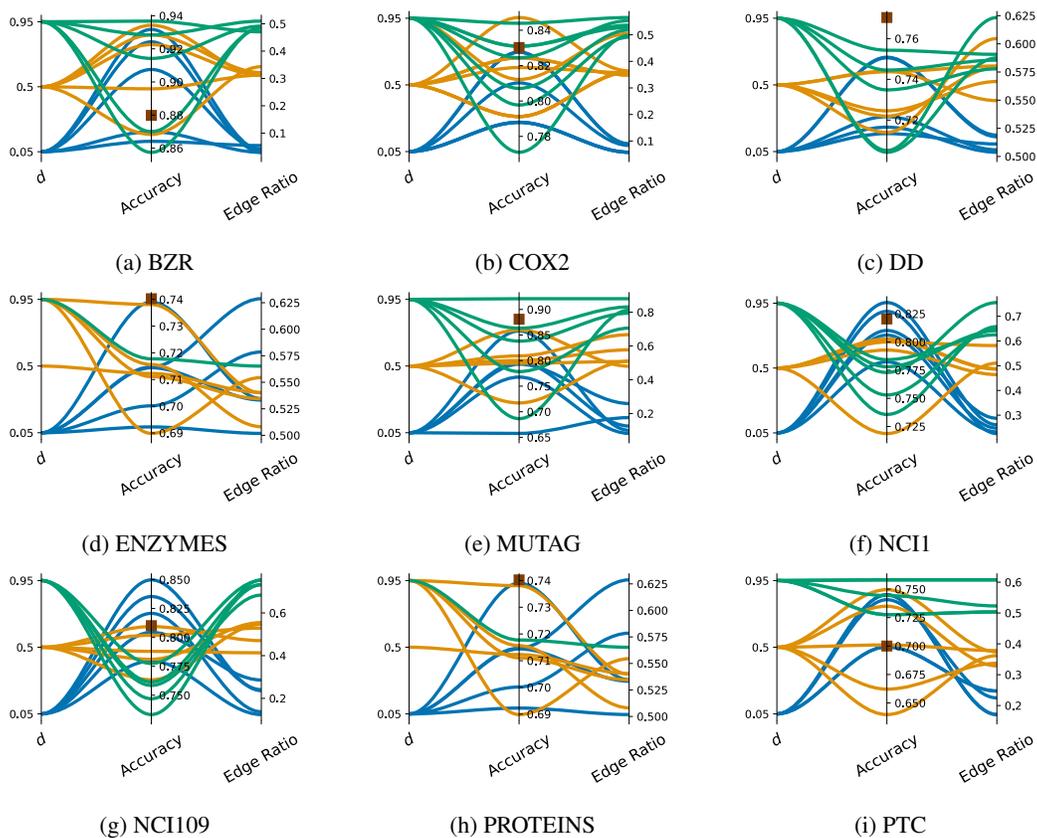


Figure 9: Ablation study on  $d$  for  $GCIP_E$ , controlling the importance given to performance or sparsity. The brown solid rectangle on the accuracy axis represents the baseline GIN (baseline) accuracy.

Table 5: Test set accuracy results. We show the mean over 5 independent runs and the standard deviation as the subindex. The last row includes the average ranking of the model across datasets. Best performing models on average are indicated in bold.

Dataset	Full Models		Sparse Models			
	GIN	GCN	GCIP <sub>GIN-N</sub>	GCIP <sub>GIN-E</sub>	GCIP <sub>GCN-N</sub>	GCIP <sub>GCN-E</sub>
BZR	81.95 <sub>2.78</sub>	<b>85.37</b> <sub>2.44</sub>	79.90 <sub>4.29</sub>	82.39 <sub>5.07</sub>	79.02 <sub>3.22</sub>	79.54 <sub>8.85</sub>
COX2	<b>84.58</b> <sub>4.56</sub>	78.75 <sub>4.52</sub>	75.49 <sub>6.33</sub>	83.59 <sub>2.82</sub>	77.81 <sub>5.03</sub>	79.05 <sub>3.48</sub>
DD	73.11 <sub>2.45</sub>	<b>76.97</b> <sub>2.19</sub>	75.51 <sub>1.86</sub>	74.40 <sub>2.12</sub>	68.35 <sub>2.75</sub>	70.02 <sub>3.92</sub>
ENZYMES	72.33 <sub>4.35</sub>	<b>72.67</b> <sub>5.48</sub>	63.70 <sub>3.89</sub>	49.53 <sub>5.29</sub>	50.88 <sub>6.88</sub>	56.02 <sub>7.60</sub>
MUTAG	<b>86.00</b> <sub>4.18</sub>	79.00 <sub>7.42</sub>	74.63 <sub>2.12</sub>	74.32 <sub>6.34</sub>	69.89 <sub>4.91</sub>	71.59 <sub>6.08</sub>
NCI1	82.04 <sub>1.26</sub>	<b>82.29</b> <sub>1.05</sub>	73.66 <sub>2.74</sub>	73.51 <sub>1.02</sub>	68.48 <sub>2.03</sub>	69.69 <sub>1.47</sub>
NCI109	<b>81.59</b> <sub>1.91</sub>	79.90 <sub>1.79</sub>	71.88 <sub>3.94</sub>	72.48 <sub>2.78</sub>	64.74 <sub>2.71</sub>	66.93 <sub>2.10</sub>
PROTEINS	72.86 <sub>4.22</sub>	71.19 <sub>2.30</sub>	73.49 <sub>7.27</sub>	<b>73.50</b> <sub>5.03</sub>	69.18 <sub>3.39</sub>	70.23 <sub>3.09</sub>
PTC	59.44 <sub>8.91</sub>	55.00 <sub>3.62</sub>	52.92 <sub>8.67</sub>	<b>66.06</b> <sub>4.31</sub>	61.14 <sub>9.79</sub>	60.00 <sub>10.41</sub>
<b>Accuracy rank</b>	<b>2.33</b>	2.33	3.67	2.89	5.33	4.44

Table 6: Test set accuracy results. We show the mean over 5 independent runs and the standard deviation as the subindex. The last row includes the average ranking of the model across datasets. Best performing models on average are indicated in bold.

Dataset	Sparse Models			
	RandomNode	RandomEdge	GCIP <sub>N</sub>	GCIP <sub>E</sub>
BZR	64.39 <sub>5.62</sub>	71.71 <sub>2.09</sub>	79.90 <sub>4.29</sub>	<b>82.39</b> <sub>5.07</sub>
COX2	70.42 <sub>8.51</sub>	78.75 <sub>9.40</sub>	75.49 <sub>6.33</sub>	<b>83.59</b> <sub>2.82</sub>
DD	39.66 <sub>7.64</sub>	41.68 <sub>3.71</sub>	<b>75.51</b> <sub>1.86</sub>	74.40 <sub>2.12</sub>
ENZYMES	27.67 <sub>6.30</sub>	29.00 <sub>7.23</sub>	<b>63.70</b> <sub>3.89</sub>	49.53 <sub>5.29</sub>
MUTAG	<b>75.00</b> <sub>7.07</sub>	69.00 <sub>17.57</sub>	74.63 <sub>2.12</sub>	74.32 <sub>6.34</sub>
NCI1	70.07 <sub>2.31</sub>	70.36 <sub>3.83</sub>	<b>73.66</b> <sub>2.74</sub>	73.51 <sub>1.02</sub>
NCI109	68.12 <sub>1.70</sub>	67.63 <sub>2.15</sub>	71.88 <sub>3.94</sub>	<b>72.48</b> <sub>2.78</sub>
PROTEINS	71.01 <sub>5.13</sub>	71.19 <sub>3.55</sub>	73.49 <sub>7.27</sub>	<b>73.50</b> <sub>5.03</sub>
PTC	-	-	52.92 <sub>8.67</sub>	<b>66.06</b> <sub>4.31</sub>
<b>Accuracy rank</b>	3.50	3.12	1.78	<b>1.56</b>

### C.3 COMPARISON WITH DIFFERENT BASE ARCHITECTURE

Here, we analyze the results of how switching the base model affects the performance. Table 5 summarizes the results for two base architectures GCN and GIN. It can be observed that even when the base architecture changes our models GCIP<sub>N</sub> and GCIP<sub>E</sub> achieve considerable accuracy compared to the respective baseline.

### C.4 COMPARISON WITH RANDOM SPARSER GRAPHS

In this section we compare GCIP with a baseline sparsity approach: we obtain a sparser graph by randomly removing a fixed ratio of nodes or edges of each graph in the dataset. We denote the resulting approaches RandomNode and RandomEdge respectively. We cross-validate the ratio of nodes and edges kept and we used GIN as the base model.

Table 6 shows the accuracy results, and Table 7 shows the results of sparsity. In Table 6, we can observe that GCIP<sub>N</sub> and GCIP<sub>E</sub> perform much better in accuracy compared to the RandomNode and RandomEdge models. Additionally, in Table 7, we show that GCIP<sub>N</sub> achieves the highest Node and Edge rank.

Table 7: Node/Edge ratio (shown in %) for the graphs in the test set for *Sparse* models on 9 different datasets. Numbers in parentheses indicate the ranking of the model for each dataset. The last two rows indicate the average ranking of the models across datasets in terms of node and edge sparsity, respectively. Best performing models on average are indicated in bold.

Dataset	Nodes/Edges %		Sparse Models				
	RandomNode	RandomEdge	TopK <sub>hard</sub>	GCIP <sub>N</sub>	GCIP <sub>E</sub>		
BZR	Node Ratio (%)	<b>18.84</b> ± 0.13 ( <b>1.00</b> )	100.00 ± 0.00 (4.00)	31.28 ± 0.04 (3.00)	18.87 ± 1.67 (2.00)	100.00 ± 0.00	
	Edge Ratio (%)	<b>2.99</b> ± 0.25 ( <b>1.00</b> )	9.45 ± 0.04 (2.00)	22.09 ± 1.86 (4.00)	12.80 ± 1.37 (3.00)	29.07 ± 2.89 (5.00)	
COX2	Node Ratio (%)	78.89 ± 0.09 (3.00)	100.00 ± 0.00 (4.00)	50.63 ± 0.10 (2.00)	<b>17.48</b> ± <b>1.72</b> ( <b>1.00</b> )	100.00 ± 0.00	
	Edge Ratio (%)	62.00 ± 1.50 (5.00)	39.58 ± 0.05 (3.00)	44.50 ± 3.68 (4.00)	<b>11.52</b> ± <b>2.82</b> ( <b>1.00</b> )	20.41 ± 1.33 (2.00)	
DD	Node Ratio (%)	<b>9.75</b> ± <b>0.02</b> ( <b>1.00</b> )	100.00 ± 0.00 (4.00)	90.25 ± 0.02 (3.00)	66.64 ± 2.73 (2.00)	100.00 ± 0.00	
	Edge Ratio (%)	<b>0.88</b> ± <b>0.03</b> ( <b>1.00</b> )	69.95 ± 0.00 (4.00)	81.19 ± 0.34 (5.00)	44.20 ± 3.40 (2.00)	58.14 ± 0.77 (3.00)	
ENZYMES	Node Ratio (%)	<b>78.34</b> ± <b>0.27</b> ( <b>1.00</b> )	100.00 ± 0.00 (4.00)	91.74 ± 0.08 (3.00)	81.84 ± 1.76 (2.00)	100.00 ± 0.00	
	Edge Ratio (%)	60.36 ± 0.52 (3.00)	<b>39.45</b> ± <b>0.14</b> ( <b>1.00</b> )	84.25 ± 0.63 (5.00)	67.16 ± 2.88 (4.00)	47.16 ± 12.87 (2.00)	
MUTAG	Node Ratio (%)	86.88 ± 0.73 (3.00)	100.00 ± 0.00 (4.00)	73.04 ± 0.89 (2.00)	<b>66.00</b> ± <b>9.87</b> ( <b>1.00</b> )	100.00 ± 0.00	
	Edge Ratio (%)	74.34 ± 2.03 (3.00)	78.95 ± 0.16 (4.00)	<b>55.00</b> ± <b>1.65</b> ( <b>1.00</b> )	62.15 ± 14.50 (2.00)	95.42 ± 6.27 (5.00)	
NCII	Node Ratio (%)	88.27 ± 0.09 (3.00)	100.00 ± 0.00 (4.00)	71.78 ± 0.05 (2.00)	<b>32.08</b> ± <b>2.77</b> ( <b>1.00</b> )	100.00 ± 0.00	
	Edge Ratio (%)	77.45 ± 0.23 (4.00)	89.25 ± 0.02 (5.00)	48.82 ± 0.52 (3.00)	<b>19.89</b> ± <b>5.17</b> ( <b>1.00</b> )	25.69 ± 2.02 (2.00)	
NCII09	Node Ratio (%)	88.22 ± 0.03 (2.00)	100.00 ± 0.00 (4.00)	91.74 ± 0.06 (3.00)	<b>53.88</b> ± <b>2.06</b> ( <b>1.00</b> )	100.00 ± 0.00	
	Edge Ratio (%)	77.54 ± 0.15 (3.00)	79.24 ± 0.04 (4.00)	89.27 ± 0.33 (5.00)	41.71 ± 4.09 (2.00)	<b>40.65</b> ± <b>5.76</b> ( <b>1.00</b> )	
PROTEINS	Node Ratio (%)	<b>67.55</b> ± <b>0.14</b> ( <b>1.00</b> )	100.00 ± 0.00 (4.00)	92.27 ± 0.24 (3.00)	85.69 ± 14.06 (2.00)	100.00 ± 0.00	
	Edge Ratio (%)	44.30 ± 0.18 (2.00)	<b>9.38</b> ± <b>0.08</b> ( <b>1.00</b> )	85.76 ± 0.41 (5.00)	75.63 ± 22.91 (4.00)	47.60 ± 8.10 (3.00)	
PTC	Node Ratio (%)	-	-	94.90 ± 1.08 (2.00)	<b>39.16</b> ± <b>7.52</b> ( <b>1.00</b> )	100.00 ± 0.00	
	Edge Ratio (%)	-	-	89.94 ± 1.95 (3.00)	<b>18.06</b> ± <b>8.06</b> ( <b>1.00</b> )	52.30 ± 1.68 (2.00)	
Average Rank	Node Rank	1.88	4.00	2.56	<b>1.44</b>	-	
	Edge Rank	2.75	3.00	3.89	<b>2.22</b>	2.78	

Table 8: The table below shows the average training time along with the standard deviation of each model on each dataset for one epoch with a batch size of one.

Dataset	Full Models		Sparse Models			
	GIN	DiffPool	SUGAR	TopK <sub>hard</sub>	GCIP <sub>N</sub>	GCIP <sub>E</sub>
BZR	2.64 <sub>0.18</sub>	4.23 <sub>0.12</sub>	-	3.44 <sub>0.27</sub>	186.96 <sub>61.70</sub>	121.94 <sub>2.00</sub>
COX2	3.07 <sub>0.15</sub>	4.83 <sub>0.17</sub>	-	10.00 <sub>2.70</sub>	411.93 <sub>85.97</sub>	203.93 <sub>1.86</sub>
DD	6.13 <sub>0.13</sub>	-	-	33.31 <sub>4.90</sub>	1848.21 <sub>1.12</sub>	1857.59 <sub>1.13</sub>
ENZYMES	2.40 <sub>0.04</sub>	6.26 <sub>0.31</sub>	132.86 <sub>3.77</sub>	6.08 <sub>0.30</sub>	233.78 <sub>6.65</sub>	287.95 <sub>2.20</sub>
MUTAG	1.09 <sub>0.03</sub>	1.63 <sub>0.04</sub>	47.82 <sub>3.45</sub>	1.71 <sub>0.19</sub>	24.97 <sub>1.55</sub>	22.94 <sub>0.30</sub>
NCI1	15.44 <sub>0.19</sub>	41.74 <sub>3.03</sub>	1143.85 <sub>45.27</sub>	48.05 <sub>8.12</sub>	1596.95 <sub>107.83</sub>	1826.41 <sub>51.06</sub>
NCI109	16.68 <sub>1.15</sub>	43.73 <sub>0.44</sub>	1195.10 <sub>47.42</sub>	43.46 <sub>3.09</sub>	1533.59 <sub>15.51</sub>	1756.84 <sub>37.12</sub>
PROTEINS	5.80 <sub>0.23</sub>	11.79 <sub>1.19</sub>	511.21 <sub>22.63</sub>	13.58 <sub>2.58</sub>	484.97 <sub>14.74</sub>	505.89 <sub>2.70</sub>
PTC	1.68 <sub>0.16</sub>	3.00 <sub>0.07</sub>	95.78 <sub>1.47</sub>	3.28 <sub>0.18</sub>	64.60 <sub>1.23</sub>	67.19 <sub>0.32</sub>

## D TIME COMPLEXITY

Here, we analyze the training and inference times of all models on each dataset. We run each model on every dataset for a single epoch, utilizing a batch size of one. To ensure statistical significance and reliability in our measurements, we repeat each experiment ten times. This repetition enables us to calculate the average training and inference times accurately. Such a meticulous approach guarantees robust and precise performance assessment for our models across the diverse set of datasets under examination.

Table 8 and Table 9 show the complete results for the training and inferences times, respectively. If we focus on Table 8, we observe the models that only rely on a single forward pass are considerably faster than the models that rely on a reinforcement learning-based approach. Still, we observe that GCIP achieves comparable speed as SUGAR, and sometimes it is faster, e.g., with PTC or MUTAG. In terms of inference time, in Table 9, we can observe that even tho GCIP is still slower than the *Full Models*, it is considerably faster than SUGAR.

Table 9: The table below shows the average inference time along with the standard deviation of each model on each dataset.

Dataset	Full Models		Sparse Models			
	GIN	DiffPool	SUGAR	TopK <sub>hard</sub>	GCIP <sub>N</sub>	GCIP <sub>E</sub>
BZR	0.0025 <sub>0.0001</sub>	0.1447 <sub>0.0048</sub>	-	0.0045 <sub>0.0003</sub>	0.0540 <sub>0.0257</sub>	0.0425 <sub>0.0011</sub>
COX2	0.0026 <sub>0.0001</sub>	0.1550 <sub>0.0093</sub>	-	0.0062 <sub>0.0008</sub>	0.1031 <sub>0.0349</sub>	0.0526 <sub>0.0014</sub>
DD	0.0028 <sub>0.0004</sub>	-	-	0.0085 <sub>0.0006</sub>	0.2262 <sub>0.0842</sub>	0.1206 <sub>0.0016</sub>
ENZYMES	0.0024 <sub>0.0000</sub>	0.1992 <sub>0.0080</sub>	0.1399 <sub>0.0094</sub>	0.0063 <sub>0.0007</sub>	0.0496 <sub>0.0014</sub>	0.0625 <sub>0.0019</sub>
MUTAG	0.0029 <sub>0.0001</sub>	0.0596 <sub>0.0034</sub>	0.0135 <sub>0.0036</sub>	0.0065 <sub>0.0017</sub>	0.0448 <sub>0.0056</sub>	0.0415 <sub>0.0012</sub>
NCI1	0.0022 <sub>0.0001</sub>	1.3756 <sub>0.1246</sub>	1.0187 <sub>0.1040</sub>	0.0058 <sub>0.0004</sub>	0.0530 <sub>0.0026</sub>	0.1301 <sub>0.0602</sub>
NCI109	0.0023 <sub>0.0001</sub>	1.3685 <sub>0.0988</sub>	0.9477 <sub>0.0132</sub>	0.0055 <sub>0.0002</sub>	0.0500 <sub>0.0006</sub>	0.0563 <sub>0.0009</sub>
PROTEINS	0.0025 <sub>0.0001</sub>	0.3835 <sub>0.0619</sub>	3.9375 <sub>0.0347</sub>	0.0058 <sub>0.0007</sub>	0.0437 <sub>0.0016</sub>	0.0439 <sub>0.0008</sub>
PTC	0.0028 <sub>0.0001</sub>	0.1076 <sub>0.0050</sub>	0.0398 <sub>0.0099</sub>	0.0058 <sub>0.0005</sub>	0.0395 <sub>0.0013</sub>	0.0418 <sub>0.0005</sub>