

Learning Network Dismantling Without Handcrafted Inputs

Haozhe Tian¹, Pietro Ferraro¹, Robert Shorten¹, Mahdi Jalili², Homayoun Hamedmoghadam^{1*}

¹Dyson School of Design Engineering, Imperial College London, United Kingdom

²School of Engineering, RMIT University, Australia

ht721@ic.ac.uk, p.ferraro@ic.ac.uk, r.shorten@ic.ac.uk, mahdi.jalili@rmit.edu.au, h.hamed@ic.ac.uk

Abstract

The application of message-passing Graph Neural Networks has been a breakthrough for important network science problems. However, the competitive performance often relies on using handcrafted structural features as inputs, which increases computational cost and introduces bias into the otherwise purely data-driven network representations. Here, we eliminate the need for handcrafted features by introducing an attention mechanism and utilizing message-iteration profiles, in addition to an effective algorithmic approach to generate a structurally diverse training set of small synthetic networks. Thereby, we build an expressive message-passing framework and use it to efficiently solve the NP-hard problem of Network Dismantling, virtually equivalent to vital node identification, with significant real-world applications. Trained solely on diversified synthetic networks, our proposed model—MIND: Message Iteration Network Dismantler—generalizes to large, unseen real networks with millions of nodes, outperforming state-of-the-art network dismantling methods. Increased efficiency and generalizability of the proposed model can be leveraged beyond dismantling in a range of complex network problems.

Code — <https://github.com/HaozheTian/MIND-ND>

Extended version — <https://arxiv.org/pdf/2508.00706>

Introduction

Network dismantling is the problem of finding the sequence of node removals that most rapidly fragments a network into isolated components (Braunstein et al. 2016; Ren et al. 2019). Finding dismantling solutions is equivalent to the identification of vital components of the network system, and has profound real-world applications, such as breaking criminal organizations by arresting the key members (Ribeiro et al. 2018), stopping epidemics with targeted vaccinations (Kitsak et al. 2010; Cohen, Havlin, and Ben-Avraham 2003), ensuring the resilience of healthcare systems via the key providers (Lo Sardo et al. 2019), and preventing wildfires by securing critical locations (Demange et al. 2025). Figure 1 visualizes network dismantling of a real-world social network (Guo, Zhang, and Yorke-Smith

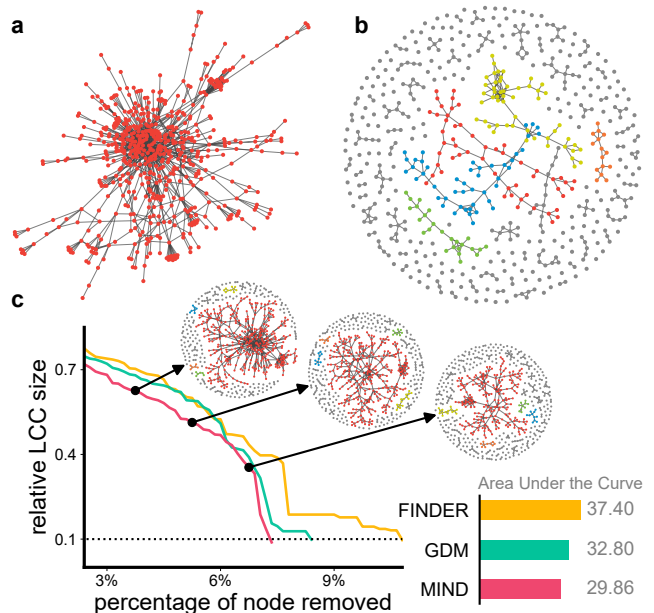


Figure 1: (a) The original social network from the FilmTrust project (Guo, Zhang, and Yorke-Smith 2016) with 610 nodes. (b) The dismantled network by MIND, down to a 10% relative Largest Connected Component (LCC) size. (c) Relative LCC size versus the fraction of nodes removed, comparing MIND with two state-of-the-art methods. (The 5 largest components are color-coded in network plots.)

2016) in action, where strategically removing a mere 7% of nodes effectively breaks it into small components.

Despite the practical significance, only approximate solutions can be sought for network dismantling, due to the NP-hard nature of the problem (Braunstein et al. 2016). Yet, the possibility of reaching a universal perception of structural roles, and the challenge of planning along the extreme breadth and depth of the search, have motivated the decades-long quest for better dismantling solutions. The early solutions use node centrality metrics as heuristics (Freeman 1977; Wandelt et al. 2018), with advancements later made by theoretical solutions to more tractable proxy problems, including

*Corresponding author: h.hamed@imperial.ac.uk.

optimal percolation (Morone and Makse 2015), graph decycling (Braunstein et al. 2016), and minimum cut (Ren et al. 2019). Recent methods use Graph Neural Networks (GNNs) to learn vector representations of nodes through iterative message-passing, which scales linearly with the number of nodes and edges and is parallelizable on GPUs (Veličković et al. 2018; Hamilton, Ying, and Leskovec 2017). The well-performing existing methods (Fan et al. 2020; Grassia, De Domenico, and Mangioni 2021) rely on handcrafted inputs to aid the inference of nodes’ importance.

We argue that using handcrafted input features i) imposes significant computational overhead, especially for large-scale networks, and ii) biases the learned dismantling strategy toward the predefined features whose effectiveness varies across network families. To address these limitations, we propose Message Iteration Network Dismantler (MIND)—a model solely based on data-driven geometric learning without manually engineered features. While it is well-established that initializing GNNs with structural heuristics improves the performance (Cui et al. 2022), MIND achieves competitive performance through pure message-passing. Specifically, MIND i) employs an expressive attention mechanism to replace the non-injective softmax normalization in existing Graph Attention Networks (GATs), and ii) leverages node embeddings from all message-passing iterations to capture crucial structural information. We demonstrate that the design of MIND enables the estimation of complex structural roles, such as those given by combinatorial centrality and spectral embedding, which are proven essential for network dismantling (Wandelt et al. 2018), and empirically show that MIND can discover new, data-driven node features that outperform known dismantling heuristics.

MIND learns to identify critical nodes and substructures with Reinforcement Learning (RL) that is trained by dismantling small synthetic networks. We introduce a training pipeline utilizing degree-preserving edge rewiring to systematically synthesize structurally diverse networks with varying levels of assortativity and modularity. Interactions with these diversified networks significantly enhance the policy’s ability to generalize to complex real-world networks. Our proposed trained policy scales well to networks with well over 1 million nodes, achieving state-of-the-art dismantling performance using only the raw incidence information, without any handcrafted input. The key contribution is the introduction of a pure geometric learning framework that can decipher the complex structural roles of network entities with surprising generalizability, resulting in the best performance to date on one of the most challenging network problems.

Related Works

Reinforcement Learning RL solves combinatorial optimization problems by learning a policy that maximizes the expected cumulative reward of action sequences in the given state space (Sutton, Barto et al. 1998; Konda and Tsitsiklis 1999). This approach enables the optimization of non-differentiable objectives via sequential decision-making; e.g., reducing network connectivity by iterative

node removals, which is infeasible through exhaustive search, and difficult for heuristic methods that do not generalize to the infinite possible network configurations. This necessitates learning from experience, where RL, especially with recent advances in Deep RL (Mnih et al. 2013; Khalil et al. 2017; Haarnoja et al. 2018), is a particularly effective solution.

Graph Neural Networks GNNs learn representations of network structures through node embeddings that are iteratively refined by aggregating the messages in each node’s neighborhood (Kipf and Welling 2016; Hamilton, Ying, and Leskovec 2017; Brody, Alon, and Yahav 2022). A shared, learnable function transforms these messages, enabling generalization to networks of arbitrary sizes. Theoretically, certain GNN architectures can distinguish almost all non-isomorphic networks even without initial node features, provided the learned function is a universal approximator (Kipf and Welling 2016; Dai, Dai, and Song 2016; Xu et al. 2019; Morris et al. 2019). However, in practice, initializing node embeddings with constant or random values often degrades the performance compared to using handcrafted features (Cui et al. 2022). The latter facilitates convergence (Oono and Suzuki 2020), but introduces bias to the learned embeddings (see the discussion on Fig. 4), as nodes with similar initial features are placed close to each other in the embedding space.

Network Dismantling via Machine Learning GNN-based embedding has significantly advanced network dismantling. Fan et al. (2020) use RL to train GNNs from experience in dismantling small random networks, but incorporate handcrafted global structural features into the GNN embeddings. Grassia, De Domenico, and Mangioni (2021) train their model on brute-force optimal dismantling sequences found for small networks, yet rely on a set of input node features (degree, neighborhood degree statistics, k -coreness, and clustering coefficient) to be calculated before being applied to dismantle a network. Khalil et al. (2017) show that without manually engineered node features, GNNs can solve other network combinatorial optimization problems, e.g., minimum vertex cover, max-cut, and the traveling-salesperson problem, yet, to the best of our knowledge, no existing method has achieved competitive dismantling performance in this setting.

Network Dismantling as an RL Problem

Let \mathcal{G} denote the universe of all possible networks and $P_{\mathcal{G}}$ be the distribution from which a network (or graph) is drawn: $G_0 = (V_0, E_0) \sim P_{\mathcal{G}}$, where V_0 is the set of nodes and E_0 is the set of edges between the nodes. At each step $t = 0, \dots, |V_0| - 1$, the network dismantling policy $\pi(v_i|G_t)$ observes $G_t = (V_t, E_t)$ and outputs a distribution over $v_i \in V_t$, from which a node is drawn $v_t \sim \pi(v_i|G_t)$ and removed from G_t (along with its incident edges), which we formulate as $G_{t+1} = G_t \setminus \{v_t\}$. With slight abuse of notation, we simplify $v_t \sim \pi(v_i|G_t)$ as $v_t = \pi(G_t)$. The standard objective for the network dismantling problem is to minimize the area under the curve (AUC) of the relative Largest Connected Component (LCC) size of the

network over the sequence of node removals, which we use to formulate policy optimization:

$$\min_{\pi} \mathbb{E} \left[\sum_{t=0}^{|V_0|-1} \frac{\text{LCC}(G_0 \setminus \{v_0, \dots, v_t\})}{|V_0|} \right], \quad (1)$$

where $\text{LCC}(\cdot)$ returns the relative size of the LCC. The optimization problem in (1) can be rewritten as the sum of rewards:

$$\max_{\pi} \mathbb{E} \left[\sum_{t=0}^{|V_0|-1} r_t \right], r_t = -\frac{\text{LCC}(G_t \setminus \{v_t\})}{|V_0|}. \quad (2)$$

Since G_{t+1} depends only on the current G_t and v_t , the problem in (2) forms a Markov Decision Process (MDP) that can be solved using RL in a data-driven manner. We also follow the standard definition of the state-action value function:

$$Q(G_t, v_i) = r_t + \mathbb{E} \left[\sum_{k=t+1}^{|V_0|-1} r_k \right], \quad (3)$$

where $Q(G_t, v_i)$ denotes the expected cumulative return (i.e., expected future AUC) starting with the removal of node v_i in network G_t and thereafter following the policy π .

Specifically, we solve (2) using an Actor-Critic RL algorithm, where the actor corresponds to the dismantling policy $\pi(v_i|G_t)$ and the critic to the state-action value function $Q(G_t, v_i)$. Each training iteration consists of two sub-processes: *value estimation* and *policy improvement*. The Actor-Critic framework features *experience replay*, which increases sample efficiency; each training iteration is performed on a randomly sampled batch of historical state transitions $\mathcal{B} \subseteq \{(G_0, v_0, r_0, G_1), \dots, (G_t, v_t, r_t, G_{t+1})\}$. In *value estimation*, MIND estimates $Q(G_t, v_i)$ with the Bellman equation:

$$Q(G_t, v_i) \approx \mathbb{E}_{\mathcal{B}} [r_t + Q(G_{t+1}, \pi(G_{t+1}))]. \quad (4)$$

Then, the *policy improvement* updates policy π by solving:

$$\hat{\pi} = \operatorname{argmax}_{\pi} \mathbb{E}_{G_t \in \mathcal{B}} [Q(G_t, \pi(G_t))]. \quad (5)$$

The batch \mathcal{B} is sampled from trajectories generated on networks from the same distribution $P_{\mathcal{G}}$ and by the same policy $\pi(v_i|G_t)$ (as in (2)), therefore, the maximization in Eq. (5) corresponds to a Monte Carlo approximation of the original objective in (2).

Methodology

To learn the representation of complex networks that is generalizable across all networks in \mathcal{G} , MIND employs a GNN-based RL framework, where both the state-action value function $Q(G_t, v_i)$ and the policy $\pi(v_i|G_t)$ are parameterized by encoder-decoder neural networks. The encoder GNNs take the adjacency representation of G_t as input and extract node embeddings z_i , which capture the structural role of each node $v_i \in V_t$. The decoders then map each z_i to a scalar score, i.e., the state-action

value of removing v_i in the $Q(G_t, v_i)$ decoder, and the probability of selecting v_i for removal in the $\pi(v_i|G_t)$ decoder. Since the encoders use a permutation-invariant GNN and the decoders are shared across all nodes, this architecture naturally handles networks of varying sizes and calculations of $Q(G_t, v_i)$ and $\pi(v_i|G_t)$ for any (G_t, v_i) pair.

GNN Encoder

To learn network representations not biased by the selection of handcrafted node features, the GNN encoder of MIND initializes each node $v_i \in V_t$ with a set of H all-ones vectors, $\{e_i^h = \mathbf{1}_F | h = 1, 2, \dots, H\}$, each e_i^h serving as a *head*, allowing for simultaneous encoding of diverse structural information. We propose a GNN encoder that incorporates two mechanisms (detailed in this section) that enable effective network representation learning with simple all-ones initialization.

All-to-One Attention Mechanism At each message-passing iteration, the embedding vector e_i^h of node v_i is updated using the following rule:

$$\hat{e}_i^h = \alpha_i^h W_{\sigma}^h e_i^h + \sum_{j \in \mathcal{N}(i)} \alpha_{i,j}^h W_{\nu}^h e_j^h, \quad (6)$$

where $W_{\sigma}^h, W_{\nu}^h \in \mathbb{R}^{F \times F}$ are learnable weight matrices. We propose the attention mechanism (MIND-AM) below to calculate the coefficients α_i^h and $\alpha_{i,j}^h$:

$$\alpha_i^h = \text{MLP}_{\sigma}^h \left(\left[\begin{array}{c} H \\ \parallel \\ W_{\sigma}^h e_i^h \end{array} \right] \right), \quad (\text{MIND-AM})$$

$$\alpha_{i,j}^h = \text{MLP}_{\nu}^h \left(\left[\begin{array}{c} H \\ \parallel \\ W_{\sigma}^h e_i^h \end{array} \right] + \left[\begin{array}{c} H \\ \parallel \\ W_{\nu}^h e_j^h \end{array} \right] \right),$$

where $\left[\begin{array}{c} H \\ \parallel \\ x^h \end{array} \right]$ is a vector concatenation as $[x^1 \parallel \dots \parallel x^H]$, and $\text{MLP}_{\sigma}^h, \text{MLP}_{\nu}^h : \mathbb{R}^{HF} \rightarrow (0, 1)$ are head-specific neural networks with sigmoid-squashed outputs. Equation (6) uses attention coefficients α^h to selectively aggregate messages from different neighbors, similar to the state-of-the-art GATs (Veličković et al. 2018; Brody, Alon, and Yahav 2022). However, GATs do not learn when $e_i^h = \mathbf{1}_F$ for all i and h , since softmax-normalization of α^h keeps node embeddings identical over message-passing iterations (as demonstrated in Fig. 5).

The idea behind MIND-AM is to employ an attention mechanism that i) eliminates the need for softmax normalization of α^h and thus preserves injectivity over the multiset $\{e_j^h : j \in \mathcal{N}(i)\}$, and ii) controls the explosion of $|e_i^h|$ without explicit normalization. Equations (MIND-AM) achieve the above by computing each head’s attention coefficient α^h using features from all heads. Thereby, our encoder automatically learns to leverage node information (e.g., local degree-like features) captured in other heads to normalize messages and prevent feature explosion.

Message Iteration Profiles Let $e_i^{(k)}$ denote the embedding vector of node v_i after the k -th message-passing iterations, calculated by concatenating the embeddings

across all heads: $e_i^{(k)} = \left[\parallel_{h=1}^H e_i^h \right]$, at layer k . MIND computes the Message Profile (MIND-MP) as the final node embedding z_i , i.e., the profile of embeddings over all message-passing iterations:

$$z_i = \text{MLP}_\zeta \left(\left[\parallel_{k=1}^K e_i^{(k)} \right] \right), \quad (\text{MIND-MP})$$

where MLP_ζ^h is a shared neural network between all nodes.

The first motivation for MIND-MP is the well-known issue of over-smoothing in node embeddings caused by iterative message-passing (Li, Han, and Wu 2018; Oono and Suzuki 2020). In Appendix A, **Theorem 1**, we show that the embeddings $e_i^{(k)}$ for all nodes $v_i \in V_t$ tend to converge to the primary eigenvector of the message-passing operator as k increases. MIND-MP retains local structural information from early iterations, thereby preserving the diversity of node embeddings.

The second motivation for MIND-MP is to extract crucial structural information that can only be obtained by jointly considering all message-passing iterations. Although $e_i^{(k)}$ converges as k increases, nodes converge at different rates, depending on their centrality (Hage and Harary 1995), as more central nodes begin aggregating information from the entire network earlier. Further theoretical insights into MIND’s expressiveness are provided in Appendix B, where we show that by learning the message-passing operator in **Lemma 1**, MIND can approximate the Fiedler vector, a widely-used spectral heuristic in network dismantling literature (Wandelt et al. 2018; Grassia, De Domenico, and Mangioni 2021).

NN Decoder

In addition to the node embeddings z_i , which encode the structural roles of individual nodes, we introduce a synthetic *omni-node* v_o to G_t . Each node $v_i \in V_t$ is connected to v_o via a directed edge, enabling one-way message-passing from all nodes to v_o . This design allows the resulting embedding z_o from the GNN Encoder to aggregate information from the entire network and represent the global state of G_t . Both the omni-node and individual node embeddings are passed to the decoders to enable state-aware decision-making. In particular, as formulated below, the Q decoder learns to estimate the remaining dismantling AUC, while the π decoder predicts the relative importance of each node for the next removal step:

$$\begin{aligned} Q(G_t, v_i) &= \text{MLP}_\theta([z_i || z_o]), \\ \pi(v_i | G_t) &= \text{MLP}_\phi([z_i || z_o]). \end{aligned} \quad (7)$$

By leveraging both local information (through z_i) and global information (through z_o), the learned network dismantling policy can perform long-term planning and adapt based on the current state of dismantling. The neural networks MLP_θ and MLP_ϕ are shared across all nodes, enabling MIND to generalize across networks of varying sizes.

Systematically Diversified Training Networks

Our goal is to train a universal dismantler that generalizes across all $G_0 \in \mathcal{G}$. So, it is essential to train on diverse

network configurations. For this purpose, the common practice is to generate networks of different sizes (and densities) using random graph models. The significance of the famous graph models, to an extent however, does not reflect their representativeness of the real (or possible) networks (and arguably has more to do with tractable mathematical properties). Here, we propose a systematic procedure to generate random training networks that better reflect the structural diversity of real-world networks. In short, the proposed procedure takes small (100-200 nodes) random networks of different degree distributions and introduces different levels of modularity and degree-assortativity by randomizing the configurations (keeping the degree sequences fixed); this also attenuates the geometrical properties inherited from the graph generation models.

We first synthesize 10,000 random networks using Linear Preferential Attachment (LPA) (Newman 2018), Copying Model (Kumar et al. 2000), and Erdos-Renyi (ER) (Erdos and Renyi 1959) models. To enhance the structural diversity, we apply degree-preserving edge rewirings to induce different types of node mixings. Specifically, we perform random edge rewirings that either favor or discourage connections between nodes with similar labels, either by degree to create varying levels of degree assortativity (assortative, uncorrelated, disassortative), or randomly to induce varying levels of modularity (modular, random, and multipartite). See Appendix C for further details on the generation process.

Entropy-Regularized Policy Learning

To train MIND for solving (2), we perform multiple dismantling episodes, each beginning with a network randomly sampled from the training set. The specific RL algorithm (detailed in Appendix D, Algorithm 1) is based on Soft Actor-Critic (SAC) (Haarnoja et al. 2018), chosen for its high sample efficiency and its ability to encourage effective exploration via entropy regularization. However, unlike the original SAC, which handles continuous action spaces via Monte Carlo sampling, the action space V_t here is discrete, allowing MIND to directly compute the expectation of the Q -value under the current policy for each G_t as:

$$\mathbb{E}_\pi [Q(G_t, v_t)] = \sum_{v_i \in V_t} \pi(v_i | G_t) Q(G_t, v_i). \quad (8)$$

Experiments

We compare the performance of MIND with a comprehensive set of baseline methods on both real-world and synthetic networks (Braunstein et al. 2016; Clusella et al. 2016; Ren et al. 2019; Fan et al. 2020; Grassia, De Domenico, and Mangioni 2021). The baselines represent both the classic methods and the state-of-the-art, identified in the recent review by Artime et al. (2024), and categorized as i) *Centrality Heuristics*: Adaptive Degree (AD), Betweenness Centrality (BC), and PageRank (PR); ii) *Approximate theory methods*: Min-Sum (MS), Explosive Immunization (EI), and Generalized Network Dismantling (GND); and iii) *Machine Learning methods*: FINDER and GDM. MIND has the lowest computational complexity

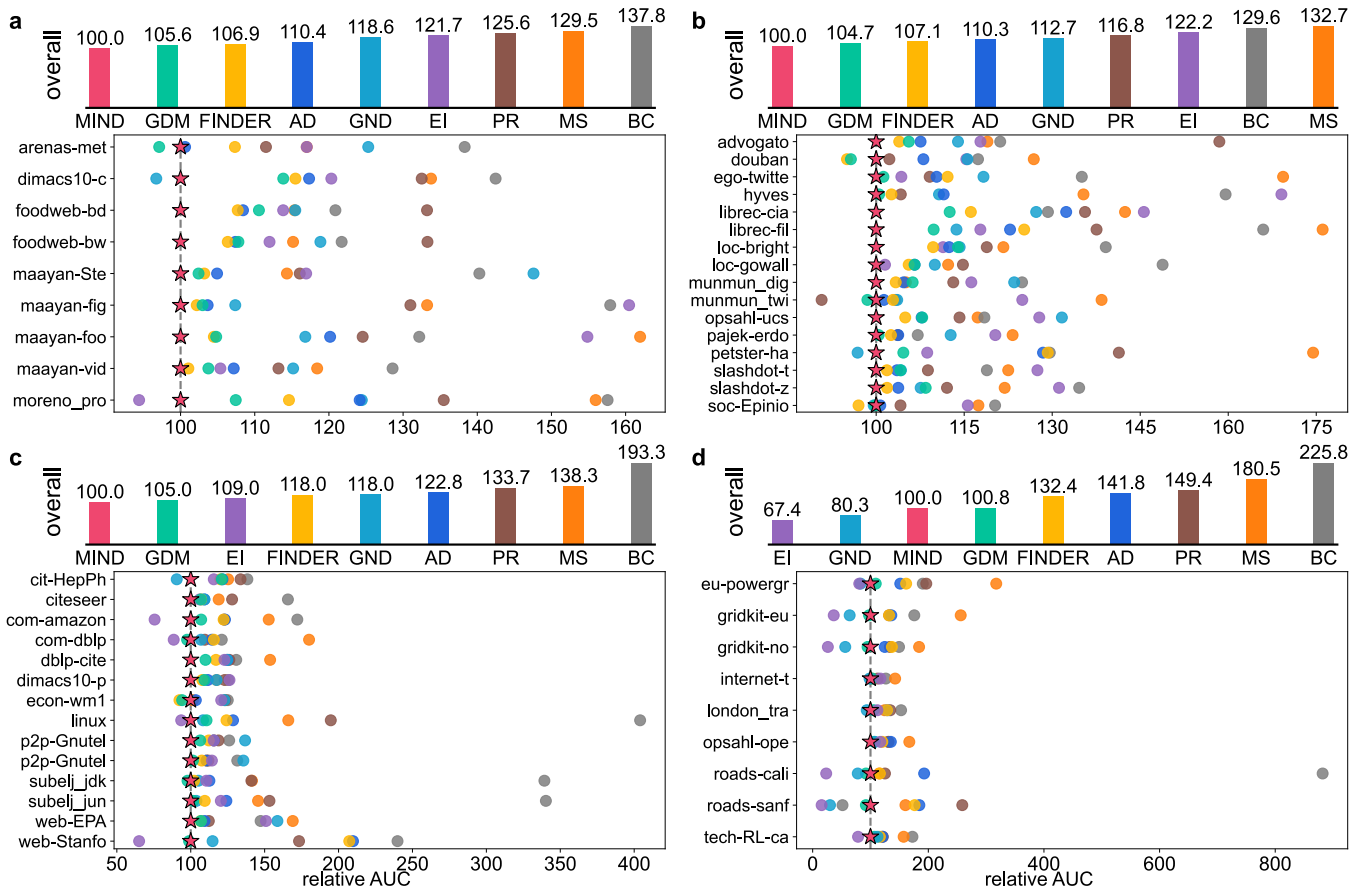


Figure 2: Dismantling performance of MIND and the baseline methods on (a) biological, (b) social, (c) information, and (d) technological networks. The scatter plots display the AUC of dismantling for all methods normalized relative to that of MIND at 100 (AUC above 100 denotes worse performance than MIND). The bar plots summarize the overall performance of the methods in each network domain, with shorter bars corresponding to lower average AUC and thus stronger dismantling performance.

among machine learning-based dismantling methods (see Table 1), as a result of not requiring the computation of handcrafted node features for embedding initialization. Compared to the approximated theory baselines, MIND is also the most computationally efficient, except for the EI method only on dense networks, where $|E|$ asymptotically grows faster than $|V| \log |V|$. All baselines are implemented following their respective references (readers may refer to the summary in Table 1 of (Artime et al. 2024)). MIND is trained over 8 million dismantling episodes, each initialized with a random selection from the training set of 10,000 small synthetic networks. The detailed training setup of MIND is provided in Appendix D.

Result on Real Networks

We evaluate MIND on real-world networks across four domains, namely, biological, social, information, and technological—covering a wide range of properties and sizes from 128 to 1.4 million nodes (summarized in Appendix E, Table 4). Figure 2 reports the AUC of the dismantling curve for all methods, normalized relative to MIND for each network; the bar plots summarize the overall

Method	Complexity
MS	$\mathcal{O}(V \log V) + \mathcal{O}(E)$
EI	$\mathcal{O}(V \log V)$
GND	$\mathcal{O}(V \log^{2+\epsilon} V)$
FINDER	$\mathcal{O}(V \log V + E)$
GDM	$\mathcal{O}(V \langle d^2 \rangle + E)$
MIND	$\mathcal{O}(V + E)$

Table 1: Computational complexity of methods assuming adjacency list representation of $G = (V, E)$. ($\langle d^2 \rangle$ is the second moment of degree.)

performances in each domain. The detailed relative AUC values are provided in Appendix F.

The top three methods, ranked by overall performance across all networks, are MIND (100.0), GDM (104.13), and EI (107.96). The results demonstrate that although other machine learning baselines take advantage of handcrafted inputs, MIND consistently achieves stronger performance across all domains. This highlights that handcrafted initial embeddings, despite boosting the GNN training, do

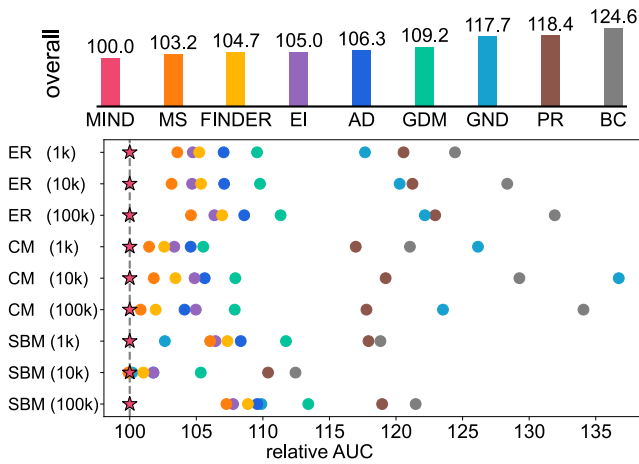


Figure 3: Dismantling performance of MIND and the baseline methods on synthetic networks (ER, CM, and SBM) of varying sizes. The scatter plot compares the dismantling performance of all methods normalized for each network relative to MIND, and the bar plot summarizes the overall performances. The AUCs are averaged over 10 realizations.




Figure 4: (a) Relative LCC size during the dismantling of an ER network with 1 k nodes. We compare the node dismantling sequence derived (by PCA) from a set of heuristics with those generated by (b) GDM and (c) MIND; Spearman rank correlation coefficient R and the regression (solid line) with confidence interval (shaded area) are shown on the plots. The heuristic removal sequence is derived from the principal component of GDM’s input node features.

not inherently yield strong dismantling performance. In contrast, MIND, empowered by our proposed MIND-AM and MIND-MP mechanisms (see the GNN Encoder section), is able to identify structurally vital nodes purely from adjacency representation, resulting in an effective network dismantling policy. In technological networks (Fig. 2d), MIND slightly underperforms EI and GND, but still outperforms all other methods, including machine learning baselines (GDM and FINDER). This can be attributed to the limited reach of GNN message-passing in technological networks with very large diameters (e.g., over 100 for gridkit-eupowergrid and gridkit-north.america).

Result on Synthetic Networks

We evaluate MIND on synthetic networks generated by the widely-adopted protocols in prior studies: (i) ER with average degree $\langle d \rangle = 4$, (ii) Configuration Model with

$\langle d \rangle = 4$ and degree distribution $P(d) \sim d^{-2.5}$, and (iii) Stochastic Block Model with group size 100, $p_{\text{intra}} = 0.1$, and $p_{\text{inter}} = 5/|V|$. From each model, we generate networks of sizes 1 k, 10 k, and 100 k, and evaluate the average AUC over 10 realizations. Figure 3 shows scatter plots of the AUC of dismantling for all methods, normalized relative to MIND for each network type, and bar plots comparing the overall performances. Note the testing networks in Fig. 3 differ from those that MIND is trained on, in both sizes and methods of generation.

From the results, we observe that MIND significantly outperforms the baselines, except for Stochastic Block Model with 1 k nodes, where the inherent community structures enable the decycling-based method (MS) to achieve a comparable performance to MIND. Notably, although GDM uses the node degree as an input feature, the learned message-passing functions extract structural information that ultimately leads to worse performance than the simple AD across all synthetic networks. Since GDM is trained on the same types of synthetic networks, this suggests that it overfits to the specific training set and loses generalizability to similar structures. This highlights the better generalizability of the RL-based (FINDER and MIND) dismantling policies compared to the supervised learning approach (GDM).

For an ER network where nodes are structurally similar, the simple AD method performs considerably better than GDM (Fig.4a). Following this observation, we investigate whether GDM’s degraded performance may be due to an inherent bias towards its handcrafted input features: degree, neighborhood degree statistics, k -coreness, and clustering coefficient. Let $X \in \mathbb{R}^{N \times 4}$ be the corresponding feature matrix for the ER network. We combine the input features by projecting X onto the principal eigenvector of $\frac{1}{N}X^T X$, and order the nodes accordingly to obtain a heuristic dismantling sequence. Figure 4b shows a significant correlation between the dismantling sequence of GDM and that of its input features. In contrast, MIND dismantling has weak to no correlation with the heuristic dismantling sequence (Fig. 4c). This corroborates that MIND gains performance by learning the underlying structural importance beyond the standard node heuristics.

Ablation Studies

GNN Design The effectiveness of the proposed MIND-AM and MIND-MP is empirically verified via ablation experiments, where we remove each design component and observe the changes in validation performance (the AUC of dismantling) during training. To calculate the validation AUC, we conduct small-scale tests and take the average performance over 20 synthetic networks generated by LPA, ER, and Watts-Strogatz (Watts and Strogatz 1998) models every 10,000 training steps. We also include the original GATv2 and GCN in the comparison, using them as the message-passing operators while keeping the rest of the MIND framework unchanged (e.g., employing message-iteration profiles instead of the final embedding). The results are shown in Fig. 5, as the mean (solid line) and standard deviation (shaded area) of validation AUC during training

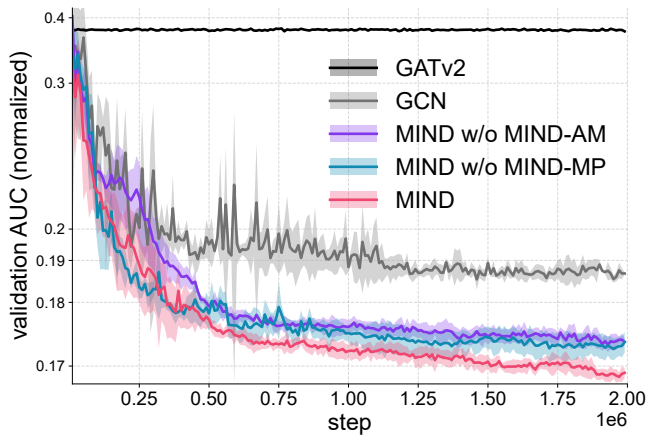


Figure 5: The validation AUC during training (mean±std). MIND is compared against: i) GATv2; ii) GCN; iii) MIND without MIND-AM (the all-to-one attention mechanism); and iv) MIND without MIND-MP (the message-profile over iterations).

over 5 independent runs.

The results demonstrate that MIND, even after removing the architectural designs proposed in this paper, still outperforms the existing GNN baselines. The original GATv2 fails to learn when the initial node embeddings are constants, due to its reliance on softmax-normalized attention coefficients. While GCN is able to learn, it achieves suboptimal performance and exhibits large fluctuations. Removing MIND-AM or MIND-MP degrades the performance of MIND. We perform t-test on the converged AUC values shown in Fig. 5 for MIND against its two ablated variants, obtaining $p = 8.1 \times 10^{-4}$ and $p = 4.5 \times 10^{-2}$, respectively, highlighting the effectiveness of the all-to-one attention mechanism to allow learnable normalization of the messages, and the benefit of utilizing information from all message-passing iterations to extract deeper structural insights.

Rewiring for Training Network Diversification To assess the effectiveness of our edge-rewiring strategy for diversifying the training networks, we compare MIND with the same model trained on the same networks, only without rewiring. The results are shown in Fig. 6, with bars depicting the effect of diversifying the training set on the performance of MIND (shorter bars correspond to higher improvement) on real networks listed in Fig. 2. For each network, the AUC of dismantling with the diversified (rewired) training set is shown as the percentage of the AUC associated with no rewiring (values below 100 indicate that training on rewired networks has led to a better dismantling policy).

The results demonstrate that rewiring the training networks yields an overall performance improvement on real networks. To analyze the results, we refer to the assortativity and modularity of the real networks in Table 4 in Appendix E. The most significant performance gains are observed for highly modular networks. For instance, in roads-california (third-to-last bar in the lower panel of




Figure 6: Performance improvement by algorithmically diversified training networks through degree-preserving edge-rewirings. Bars show the AUC of the dismantling of MIND trained on rewired networks, normalized to the baseline without rewiring (gray dotted lines).

Fig. 6), which has a modularity of 0.975, training with rewired networks led to an over 80% reduction in the AUC of dismantling. This improvement is likely linked to the rewiring-induced modularity in the otherwise non-modular synthetic networks. We also observe notable performance gains for networks with strong disassortativity. For example, in munmun_twitter_social (seventh-to-last bar in the upper panel of Fig.6), which has a degree assortativity of -0.878 , the AUC is reduced by 40%. Although the LPA and Copying Model naturally produce slightly disassortative networks, our diversifying rewirings enable the model to interact with a much wider range of (dis)assortative mixings and thereby significantly enhance the learned embedding and policy by increased exposure to different topologies.

Conclusion

Eliminating the need for initializing GNNs with handcrafted features is highly sought after. Besides dropping the feature computation overhead, featureless initialization eliminates the risk of embedding bias and grants autonomy for learning more complex embeddings, which is key to finding better solutions to the downstream problems. We tackled this with two ideas: i) building an expressive message-passing framework, and ii) exposing the model to interactions with systematically diversified network geometries, facilitating the learning of complex structural roles. The proposed model, applied to the network dismantling problem, achieved state-of-the-art performance on a comprehensive testbed of real-world networks. Note that MIND is computationally more efficient than the well-performing methods of its category, machine learning methods, as well as the well-established dismantling methods in the literature (except for EI only on dense networks). An intriguing conclusion is that the contributions of this manuscript are applicable to many important network/graph problems where an unbiased GNN embedding can be learned on synthesized diverse data and lead to breakthrough solutions.

Acknowledgments

RS and HH acknowledge the Australian Research Council Discovery Project No. DP240102585, and the support from the IOTA Foundation. RS and HT acknowledge the funding by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee (grant number 101084642). HT thanks Chi-Bach Pham for his help with some of the experiments.

References

- Artime, O.; Grassia, M.; De Domenico, M.; Gleeson, J. P.; Makse, H. A.; Mangioni, G.; Perc, M.; and Radicchi, F. 2024. Robustness and resilience of complex networks. *Nature Reviews Physics*, 6(2): 114–131.
- Braunstein, A.; Dall'Asta, L.; Semerjian, G.; and Zdeborová, L. 2016. Network dismantling. *Proceedings of the National Academy of Sciences*, 113(44): 12368–12373.
- Brody, S.; Alon, U.; and Yahav, E. 2022. How Attentive are Graph Attention Networks? In *International Conference on Learning Representations*.
- Clusella, P.; Grassberger, P.; Pérez-Reche, F. J.; and Politi, A. 2016. Immunization and targeted destruction of networks using explosive percolation. *Physical Review Letters*, 117(20): 208301.
- Cohen, R.; Havlin, S.; and Ben-Avraham, D. 2003. Efficient immunization strategies for computer networks and populations. *Physical review letters*, 91(24): 247901.
- Cui, H.; Lu, Z.; Li, P.; and Yang, C. 2022. On positional and structural node features for graph neural networks on non-attributed graphs. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 3898–3902.
- Dai, H.; Dai, B.; and Song, L. 2016. Discriminative embeddings of latent variable models for structured data. In *International Conference on Machine Learning*, 2702–2711. PMLR.
- Demange, M.; Di Fonso, A.; Di Stefano, G.; and Vittorini, P. 2025. Instantiating a Diffusion Network Model to Support Wildfire Management. *IEEE Transactions on Network Science and Engineering*.
- Erdos, P.; and Renyi, A. 1959. On Random Graphs I. *Publicationes Mathematicae Debrecen*, 6: 290–297.
- Fan, C.; Zeng, L.; Sun, Y.; and Liu, Y.-Y. 2020. Finding key players in complex networks through deep reinforcement learning. *Nature Machine Intelligence*, 2(6): 317–324.
- Freeman, L. C. 1977. A set of measures of centrality based on betweenness. *Sociometry*, 35–41.
- Grassia, M.; De Domenico, M.; and Mangioni, G. 2021. Machine learning dismantling and early-warning signals of disintegration in complex systems. *Nature Communications*, 12(1): 5190.
- Guo, G.; Zhang, J.; and Yorke-Smith, N. 2016. A novel evidence-based Bayesian similarity measure for recommender systems. *ACM Transactions on the Web*, 10(2): 1–30.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 1861–1870. Pmlr.
- Hage, P.; and Harary, F. 1995. Eccentricity and centrality in networks. *Social Networks*, 17(1): 57–63.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30.
- Khalil, E.; Dai, H.; Zhang, Y.; Dilkina, B.; and Song, L. 2017. Learning combinatorial optimization algorithms over graphs. *Advances in Neural Information Processing Systems*, 30.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Kitsak, M.; Gallos, L. K.; Havlin, S.; Liljeros, F.; Muchnik, L.; Stanley, H. E.; and Makse, H. A. 2010. Identification of influential spreaders in complex networks. *Nature Physics*, 6(11): 888–893.
- Konda, V.; and Tsitsiklis, J. 1999. Actor-critic algorithms. *Advances in Neural Information Processing Systems*, 12.
- Kumar, R.; Raghavan, P.; Rajagopalan, S.; Sivakumar, D.; Tomkins, A.; and Upfal, E. 2000. Stochastic models for the web graph. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, 57–65. IEEE.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Lo Sardo, D. R.; Thurner, S.; Sorger, J.; Duftschmid, G.; Endel, G.; and Klimek, P. 2019. Quantification of the resilience of primary care networks by stress testing the health care system. *Proceedings of the National Academy of Sciences*, 116(48): 23930–23935.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Morone, F.; and Makse, H. A. 2015. Influence maximization in complex networks through optimal percolation. *Nature*, 524(7563): 65–68.
- Morris, C.; Ritzert, M.; Fey, M.; Hamilton, W. L.; Lenssen, J. E.; Rattan, G.; and Grohe, M. 2019. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 4602–4609.
- Newman, M. 2018. *Networks*. Oxford university press.
- Oono, K.; and Suzuki, T. 2020. Graph Neural Networks Exponentially Lose Expressive Power for Node Classification. In *International Conference on Learning Representations*.
- Ren, X.-L.; Gleinig, N.; Helbing, D.; and Antulov-Fantulin, N. 2019. Generalized network dismantling. *Proceedings of the National Academy of Sciences*, 116(14): 6554–6559.

Ribeiro, H. V.; Alves, L. G.; Martins, A. F.; Lenzi, E. K.; and Perc, M. 2018. The dynamical structure of political corruption networks. *Journal of Complex Networks*, 6(6): 989–1003.

Sutton, R. S.; Barto, A. G.; et al. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. *International Conference on Learning Representations*.

Wandelt, S.; Sun, X.; Feng, D.; Zanin, M.; and Havlin, S. 2018. A comparative analysis of approaches to network-dismantling. *Scientific Reports*, 8(1): 13513.

Watts, D. J.; and Strogatz, S. H. 1998. Collective dynamics of 'small-world' networks. *Nature*, 393(6684): 440–442.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.