DDIN: Reinforcement Learning with Asymmetric GNNs for Dismantling Directed Interdependent Networks (Student Abstract)

Soumyajit Dev, Malay Bhattacharyya
Indian Statistical Institute, Kolkata
soumyajit1998dev@gmail.com, malaybhattacharyya@isical.ac.in

Abstract

Dismantling interdependent directed networks to obtain the largest mutually strongly connected component (MSCC) is an NP-hard problem. To address this, we propose a novel method, Disassembling Directed Interdependent Networks (DDIN), by synergizing Reinforcement Learning (RL) and Graph Neural Networks (GNN). We introduce asymmetric GNNs to capture the asymmetry of in/out-degree and multi-relational attention to model directed inter-layer dependencies, integrated with prioritized RL for efficient node selection in large action spaces. Our contributions include (i) a directed GraphSAGE encoder separating in/out aggregations for asymmetry, (ii) multi-relational attention fusing layer semantics, and (iii) sum-tree prioritized n-step Deep Q-Network (DQN) for efficient policy search. DDIN is evaluated on 5 directed multiplexes from biological, social, and economic domains, achieving 16-23% lower AUDC compared to known baseline heuristics.

Introduction

Interdependent directed networks model critical real-life systems (e.g., biological and cyber-physical systems, supply chains, and social media) (Barabási, Gulbahce, and Loscalzo 2011; Watts and Strogatz 1998), where directed intra-layer edges capture unidirectional flows, and undirected inter-layer dependencies symbolize mutual reinforcement. This makes them vulnerable to asymmetric propagation of failures (Vespignani 2010). Finding minimal nodes to obtain the largest mutually strongly connected component (MSCC) is NP-hard (Braunstein et al. 2016). Traditional methods like High Degree Attack (HDA) or Collective Influence (CI) work in undirected settings (Morone and Makse 2015), but fail in directed cases, leading to suboptimal dismantling (Ma, Wang, and Li 2022; Artime et al. 2024). Recently, GNNs and RL are used for the dismantling of undirected networks (Rahmede et al. 2018; Tian et al. 2025; Gu et al. 2025). Previous RL models work well on undirected networks but not for directed asymmetries (Fan et al. 2020; Grassia, De Domenico, and Mangioni 2021). We introduce DDIN, a framework using asymmetric GNNs and prioritized multi-step RL for dismantling directed interdependent networks. While the prior work handles undirected graphs via symmetric aggregations (Zhang and Wang 2022), DDIN extends them with asymmetric GraphSAGE encoders that differentiate incoming/outgoing edges, enabling zero-shot generalization to directed interdependencies. DDIN not only reduces the critical fraction of nodes for dismantling but reveals patterns in directed cascades, e.g., targeting high-out-degree hubs in one layer to disrupt cross-layer flows.

Problem Statement

In directed multiplex networks, the objective is to find the smallest node set S that reduces the size of the largest MSCC below a threshold, while accounting for directed percolation and interdependent cascading failures. Formally, consider a directed multilayer graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L})$, where $\mathcal{V}$ is the shared node set, $\mathcal{E} = \bigcup_{\ell \in \mathcal{L}} \mathcal{E}_\ell$ with $\mathcal{E}_\ell \subseteq \mathcal{V} \times \mathcal{V}$ denotes directed arcs in layer ℓ, and inter-layer dependencies are bidirectional (Bianconi 2018). The MSCC is the largest subset $C \subseteq \mathcal{V}$ where for all $u, v \in C$, there exists directed paths $u \rightarrow v$ and $v \rightarrow u$, considering paths within and across layers via interdependencies (Osat, Faqeeh, and Radicchi 2017). The dismantling cost function $F(S)$ measures the normalized MSCC size after removing S: $F(S) = \frac{\max\{|C|:C \in \text{MSCCs}(\mathcal{G} \setminus S)\}}{N}$, where $N = |\mathcal{V}|$ and MSCCs$(\cdot)$ is the set of MSCCs in the residual graph. In interdependent settings, node removal in one layer triggers cascades: if node v in layer ℓ fails, dependent nodes in other layers ℓ' fail with certain probability $p_{v,\ell'}$ (Buldyrev et al. 2010). The corresponding optimization problem is $\min_{S \subseteq \mathcal{V}} |S|$ s.t. $F(S) < \theta$, with $\theta \ll 1$ (e.g., 0.1) (Ren et al. 2019). To evaluate uniform efficiency, we define the **Average Uniform Dismantling Cost** (AUDC) $= \frac{1}{F(\mathcal{V})} \sum_{a=1}^{F(\mathcal{V})} f(S_a)/|S_a|$, where $f(S_a)$ is the cost of the removal sequence S_a, aggregated over all orders. In directed settings, $f(S_a)$ incorporates asymmetry due to in- and out-paths, revealing patterns like targeting high-out-degree hubs to disrupt cross-layer flows (Clusella et al. 2016).

Proposed Method

We generate directed interdependent networks using an enhanced Geometric Multiplex Model (GMM) (Kleineberg et al. 2017), with separate power-law exponents for in/out-degrees and directional bias α. The connection probability is $p_{ij} \propto \exp(\alpha(\kappa_i^{\text{out}} - \kappa_j^{\text{in}}))$, where κ denotes hidden degree

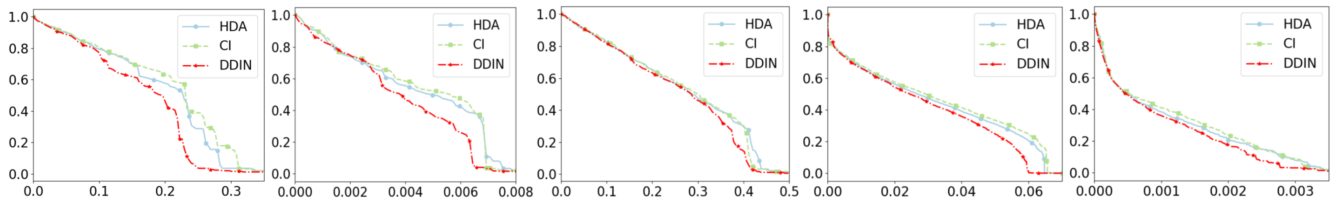

Figure 1: Dismantling curves showing normalized MSCC of residual network versus fraction of node removal cost (ϕ) for the datasets (A) C. elegans connectome (279 nodes, 3 layers), (B) Drosophila Melanogaster (8215 nodes, 7 layers), (C) FAO Trade (214 nodes, 364 layers), (D) Homo Genetic (18k nodes, 7 layers), and (E) Sanremo 2016 (56k nodes, 3 layers).

variables sampled from conditional power-law distributions. This setup embeds nodes in hyperbolic space, naturally producing scale-free directed structures that mimic asymmetric real-world interactions (Kleineberg et al. 2016). We model directed interdependent networks as a multiplex with directed intra-layer and undirected inter-layer edges, allowing asymmetric failure propagation (failure A$\rightarrow$B $\not\Rightarrow$ B$\rightarrow$A).

Encoder: It is a directed GraphSAGE model designed to extract asymmetric node embeddings. For each layer ℓ, the initial feature vector $\mathbf{x}_v$ includes out-/in-degrees and hyperbolic coordinates. The embedding update is: $\mathbf{h}_v^\ell = \text{ReLU}\Big(W_{\text{out}}\left[\mathbf{x}_v \,;\, \text{MeanAgg}\big(\{\mathbf{h}_u^\ell : u \to v\}\big)\right] + W_{\text{in}}\left[\mathbf{x}_v \,;\, \text{MaxAgg}\big(\{\mathbf{h}_u^\ell : v \to u\}\big)\right]\Big)$. Here MeanAgg$(\cdot)$ = $\frac{1}{|\mathcal{N}|}\sum_{u\in\mathcal{N}} \mathbf{h}_u^\ell$, $\mathcal{N}_{\text{out}} = \{u : u \to v\}$ and $\mathcal{N}_{\text{in}} = \{u : v \to u\}$ are out- and in-neighborhoods. The matrices $W_{\text{out}}, W_{\text{in}} \in \mathbb{R}^{d\times 2d}$ are learnable, enabling the capture of directional flows. GraphSAGE assumes symmetry (Hamilton, Ying, and Leskovec 2017), but directionality requires separating in/out neighbors (Bojchevski et al. 2020). Our asymmetric encoder aggregates incoming features with mean pooling (capturing influences) and outgoing with max pooling (capturing controls), preserving flow asymmetries.

Multi-relational attention layer: The attention coeff., $\alpha_{\ell\ell'} = \text{softmax}\left(\text{LeakyReLU}\left(\mathbf{a}^\top\left[\mathbf{W}\mathbf{h}_v^\ell \| \mathbf{W}\mathbf{h}_v^{\ell'}\right]\right)\right)$, yielding $\tilde{\mathbf{h}}_v = \sum_{\ell'\neq\ell} \alpha_{\ell\ell'}\mathbf{h}_v^{\ell'}$ (Veličković et al. 2017). Inter-layer dependencies are relational (e.g., one layer's out-edges affect another's in-edges). Multi-relational attention weights these asymmetrically, fusing semantics to capture cross-layer cascades (Hu et al. 2020), unlike uniform fusion in undirected methods (Melton and Krishnan 2023).

Decoder: It reconstructs the state of the directed network from fused embeddings $\tilde{\mathbf{h}}_v$. To predict the edge between i and j in layer ℓ, it computes $\hat{p}_{ij}^\ell = \sigma\left(\tilde{\mathbf{h}}_i^\ell W_{\text{dec}} \tilde{\mathbf{h}}_j^{\ell\top} + b^\ell\right)$, where σ is sigmoid, $W_{\text{dec}} \in \mathbb{R}^{d\times d}$ is a bilinear decoder weight, and b^ℓ is a layer bias incorporating directional α.

Loss Function and Dismantling Process: The loss combines reconstruction and RL objectives $\mathcal{L} = \mathcal{L}_{\text{RL}} + \lambda\mathcal{L}_{\text{rec}}$, where $\mathcal{L}_{\text{rec}} = -\sum_{(i,j)\in\mathcal{E}} \log \hat{p}_{ij} - \sum_{(i,j)\notin\mathcal{E}} \log(1-\hat{p}_{ij})$ is binary cross-entropy, and $\mathcal{L}_{\text{RL}}$ is the Huber TD loss $\mathcal{L}_{\text{RL}} = \frac{1}{B}\sum_i \delta(e_i)\cdot\rho_i$, with $e_i = r_i + \gamma^n \max_{a'} Q(s_{i+n}, a'; \theta^-) - Q(s_i, a_i; \theta)$, and ρ_i is the importance weight. The dismantling process uses the trained Q-network for greedy node selection. In step t, choose $a_t = \arg\max_a Q(s_t, a; \theta)$, update the network via directed MSCC recomputation (iteratively via Tarjan's algorithm (Tarjan 1972)), and continue until $F(s_t) < \theta$. The prioritized experience replay focuses on high-error samples (Schaul et al. 2015), efficient for sparse rewards in large directed graphs. The loop optimizes for minimal removals of nodes to reduce the MSCC, with rewards based on fragmentation gain.

Experimental Results

DDIN is evaluated on 5 real-world networks from different domains (De Domenico 2025). We train the model on NVIDIA A100 Tensor Core GPU with batch size = 32, $\epsilon = 0.01$, Learning Rate = 10^{-4}, $\gamma = 1$, MSCC via iterative Tarjan's algorithm. The measures AUDC and ϕ^* are calculated by averaging over 5 runs. The baselines are Directed CI (DCI) and HDA (Albert, Jeong, and Barabási 2000) (extended to multilayers). It outperforms baselines via asymmetric targeting (Table 1 shows lower AUDC, Fig. 1 shows steeper MSCC declines). DDIN enables nuanced directed cascade representation and efficient exploration.

Dataset	DCI	HDA	DDIN	Improvement
Celegans	0.2043	0.2000	**0.1621**	23.40%
Drosophila	0.0674	0.0650	**0.0561**	16.02%
FAO Trade	0.1890	0.1858	**0.1541**	20.56%
Homo Genetic	0.0149	0.0140	**0.0116**	20.53%
Sanremo 2016	0.0003	0.0003	**0.0002**	20.51%

Table 1: The AUDC values (lower is better) and relative improvement of DDIN over the state-of-the-art methods.

Conclusion

Our key contributions include directed GraphSAGE, multi-relational attention, and sum-tree prioritized n-step DQN for efficiency, scalability and adaptability. DDIN reduces the critical fraction of nodes for dismantling and reveals patterns (e.g., targeting high-out-degree hubs in one layer to disrupt cross-layer flows) in directed cascades. This provides reusable insights for cascade control, such as guiding strategies for disease containment in multilayer social networks and delaying cascading failures in interdependent critical infrastructures. We plan to handle dynamic evolving networks in future where edges change over time, incorporating partial observability for scenarios with incomplete inter-layer knowledge, and integrating hybrid optimizers.

References

Albert, R.; Jeong, H.; and Barabási, A.-L. 2000. Error and attack tolerance of complex networks. *nature*, 406(6794): 378–382.

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.

Barabási, A.-L.; Gulbahce, N.; and Loscalzo, J. 2011. Network medicine: a network-based approach to human disease. *Nature reviews genetics*, 12(1): 56–68.

Bianconi, G. 2018. *Multilayer networks: structure and function*. Oxford university press.

Bojchevski, A.; Gasteiger, J.; Perozzi, B.; Kapoor, A.; Blais, M.; Rózemberczki, B.; Lukasik, M.; and Günnemann, S. 2020. Scaling graph neural networks with approximate pagerank. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2464–2473.

Braunstein, A.; Dall'Asta, L.; Semerjian, G.; and Zdeborová, L. 2016. Network dismantling. *Proceedings of the National Academy of Sciences*, 113(44): 12368–12373.

Buldyrev, S. V.; Parshani, R.; Paul, G.; Stanley, H. E.; and Havlin, S. 2010. Catastrophic cascade of failures in interdependent networks. *Nature*, 464(7291): 1025–1028.

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.

De Domenico, M. 2025. Multilayer Network Datasets Released for Reproducibility. Personal homepage. Accessed: 2025-11-10.

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.

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.

Gu, W.; Yang, C.; Li, L.; Hou, J.; and Radicchi, F. 2025. Deep-learning-aided dismantling of interdependent networks. *Nature Machine Intelligence*, 1–12.

Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.

Hu, Z.; Dong, Y.; Wang, K.; and Sun, Y. 2020. Heterogeneous graph transformer. In *Proceedings of the web conference 2020*, 2704–2710.

Kleineberg, K.-K.; Boguná, M.; Ángeles Serrano, M.; and Papadopoulos, F. 2016. Hidden geometric correlations in real multiplex networks. *Nature Physics*, 12(11): 1076–1081.

Kleineberg, K.-K.; Buzna, L.; Papadopoulos, F.; Boguñá, M.; and Serrano, M. Á. 2017. Geometric correlations mitigate the extreme vulnerability of multiplex networks against targeted attacks. *Physical review letters*, 118(21): 218301.

Ma, J.; Wang, P.; and Li, H. 2022. Directed Network Disassembly Method Based on Non-Backtracking Matrix. *Applied Sciences*, 12(23): 12047.

Melton, J.; and Krishnan, S. 2023. muxGNN: Multiplex graph neural network for heterogeneous graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9): 11067–11078.

Morone, F.; and Makse, H. A. 2015. Influence maximization in complex networks through optimal percolation. *Nature*, 524(7563): 65–68.

Osat, S.; Faqeeh, A.; and Radicchi, F. 2017. Optimal percolation on multiplex networks. *Nature communications*, 8(1): 1540.

Rahmede, C.; Iacovacci, J.; Arenas, A.; and Bianconi, G. 2018. Centralities of nodes and influences of layers in large multiplex networks. *Journal of Complex Networks*, 6(5): 733–752.

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.

Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.

Tarjan, R. 1972. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2): 146–160.

Tian, H.; Ferraro, P.; Shorten, R.; Jalili, M.; and Hamedmoghadam, H. 2025. Learning Network Dismantling without Handcrafted Inputs. *arXiv preprint arXiv:2508.00706*.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.

Vespignani, A. 2010. The fragility of interdependency. *Nature*, 464(7291): 984–985.

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

Zhang, J.; and Wang, B. 2022. Dismantling complex networks by a neural model trained from tiny networks. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2559–2568.