Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Goal-directed graph construction using reinforcement learning

(Article begins on next page)

29 June 2024

Goal-directed graph construction using reinforcement learning

Victor-Alexandru Darvariu[1,2], Stephen Hailes[1] and Mirco Musolesi[1,2,3]

[1]Department of Computer Science, University College London, London, UK

[2]The Alan Turing Institute, London, UK

[3]Department of Computer Science and Engineering, University of Bologna, Bologna, Italy

{v.darvariu, s.hailes, m.musolesi@ucl.ac.uk}

Abstract

Graphs can be used to represent and reason about systems and a variety of metrics have been devised to quantify their global characteristics. However, little is currently known about how to construct a graph or improve an existing one given a target objective. In this work, we formulate the construction of a graph as a decision-making process in which a central agent creates topologies by trial and error and receives rewards proportional to the value of the target objective. By means of this conceptual framework, we propose an algorithm based on reinforcement learning and graph neural networks to learn graph construction and improvement strategies. Our core case study focuses on robustness to failures and attacks, a property relevant for the infrastructure and communication networks that power modern society. Experiments on synthetic and real-world graphs show that this approach can outperform existing methods while being cheaper to evaluate. It also allows generalization to out-of-sample graphs, as well as to larger out-of-distribution graphs in some cases. The approach is applicable to the optimization of other global structural properties of graphs.

1 Introduction

Graphs are mathematical abstractions that can be used to model a variety of systems, from infrastructure and biological networks to social structures. Various methods for analysing networks have been developed: these have been often used for understanding the systems themselves and range from mathematical models of how families of graphs are generated [60, 4] to measures of centrality for capturing the roles of vertices [10] and global network characteristics [44], to name but a few.

A measure that has attracted significant interest from researchers and practitioners is *robustness* [43] (sometimes called *resilience*), which is typically defined as the capacity of the graph to withstand random failures, targeted attacks on key nodes, or some combination thereof. A network is considered robust if a large fraction (*critical fraction*) of nodes have to be removed before it becomes disconnected [16], its diameter increases [2], or its largest connected component diminishes in size [9]. Previous work has focused on the robustness of communication networks such as the Internet [17] and infrastructure networks used

Figure 1: Illustration of a Graph Construction MDP (GC-MDP) trajectory. The agent is provided with a start state $S_0 = (G_0, \varnothing)$. It must make $L = 3$ edge additions over a sequence of 6 node selections (actions A_t), receiving rewards R_t proportional to the value of an objective function $\mathcal{F}$ applied to the graph. In this case, $\mathcal{F}$ quantifies the robustness of the network to targeted node removal, computed by removing nodes in decreasing order of their degree and in decreasing order of the labels if two nodes have the same degree. We observe an improvement of the robustness of the graph from $\mathcal{F}(G_0) = 0.1$ to $\mathcal{F}(G_6) = 0.7$. Actions and the corresponding edges are highlighted.

for transportation and energy distribution [14], for which resilience is a key property.

In many practical cases, an initial network is given and the only way of improving its robustness is through the modification of its structure. This problem was first approached by considering edge addition or rewiring, based on random and preferential (w.r.t. node degree) modifications [9]. Alternatively, a strategy has been proposed that uses a "greedy" modification scheme based on random edge selection and swapping if the resilience metric improves [48]. Another line of work focuses on the spectral decomposition of the graph Laplacian, and using properties such as the algebraic connectivity [57] and effective graph resistance [58] to guide modifications. While simple and interpretable, these strategies may not yield the best solutions or generalise across networks with varying characteristics and sizes. Certainly, better solutions may be found by exhaustive search, but the time complexity of exploring all the possible topologies and the cost of computing the metric render this strategy infeasible. With the goal of discovering better strategies than existing methods, we ask whether *generalisable network construction strategies for improving robustness can be learned.*

Starting from this motivation, we formalise the process of graph construction and improvement as a Markov Decision Process (MDP) in which rewards are proportional to the value of a graph-level objective function. We consider two objective functions that quantify robustness as the critical fraction of the network in the presence of random failures and targeted attacks. Inspired by recent successes of RL in solving combinatorial optimisation problems on graphs [6, 32], we make use of Graph Neural Network (GNN) architectures [26] together with the Deep Q-Network (DQN) [41] algorithm. Recent work in goal-directed graph generation and improvement considers performing edge additions for adversarially attacking GNN classifiers [19] and generating molecules with certain desirable properties using domain-specific rewards [64]. In contrast, to the best of our knowledge, this is the first time that RL is used *to learn how to construct a graph such as to optimise a global structural property.* While in this paper we focus on robustness, other intrinsic global properties of graphs, such

2

as efficiency [34] or communicability [22], could be used as optimisation targets.

The contribution of this paper is twofold. Firstly, we propose a framework for improving global structural properties of graphs, by introducing the Graph Construction Markov Decision Process (GC-MDP). Secondly, focusing on the robustness of graphs under failures and attacks as a core case study, we offer an in-depth empirical evaluation that demonstrates significant advantages over existing approaches in this domain, both in terms of the quality of the solutions found as well as the time complexity of model evaluation. Since this approach addresses the problem of building robust networks with a DQN, we name it *RNet–DQN*.

The remainder of the paper is structured as follows. We provide the definitions of the GC-MDP and the robustness measures in Section 2. Section 3 describes state and action representations for deep RL using GNNs. We present our experimental setup in Section 4, and discuss our main results in Section 5. In Section 6 we review and compare the key works in this area. Finally, we conclude and offer a discussion of avenues for future work in Section 7.

2 Robust Graph Construction as a Decision Making Problem

MDP Preliminaries. An MDP is one possible formalisation of a decision making process. The decision maker, called an *agent*, interacts with an *environment*. When in a *state* $s \in \mathcal{S}$, the agent must take an *action* a out of the set $\mathcal{A}(s)$ of valid ones, receiving a *reward* r governed by the reward function $\mathcal{R}(s, a)$. Finally, the agent finds itself in a new state s', depending on a transition model $\mathcal{P}$ that governs the joint probability distribution $P(s', a, s)$ of transitioning to state s' after taking action a in state s. This sequence of interactions gives rise to a *trajectory*. The agent's goal is to maximise the expected (possibly discounted) sum of rewards it receives over all the trajectories. The tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ defines this MDP, where $\gamma \in [0, 1]$ is a discount factor. We also define a *policy* $\pi(a|s)$, i.e., a distribution of actions over states, which fully determines the behaviour of the agent. Given a policy π, the *action-value function* $Q_\pi(s, a)$ is defined as the expected return when starting from s, taking action a, and subsequently following policy π.

Modelling Graph Construction. Let $\mathbf{G}^{(N)}$ be the set of labelled, undirected, unweighted graphs with N nodes; each such graph $G = (V, E)$ consists of a vertex set V and edge set E. Let $\mathbf{G}^{(N,m)}$ be the subset of $\mathbf{G}^{(N)}$ with $|E| = m$. We also let $\mathcal{F} \colon \mathbf{G}^{(N)} \to [0, 1]$ be an objective function, and $L \in \mathbb{N}$ be a modification budget. Given an initial graph $G_0 = (V, E_0) \in \mathbf{G}^{(N,m_0)}$, the aim is to perform a series of L edge additions to G_0 such that the resulting graph $G_* = (V, E_*)$ satisfies:

$$G_* = \operatorname*{argmax}_{G' \in \mathbf{G}'} \mathcal{F}(G'),$$
$$\text{where} \quad \mathbf{G}' = \{G = (V, E) \in \mathbf{G}^{(N,m_0+L)} \mid E_0 \subset E\}.$$

This combinatorial optimisation problem can be cast as a sequential decision-making process. In order to enable scaling to large graphs, the agent has to

3

select a node at each step, and an edge is added to the graph after every two decisions [19]. Tasks are episodic; each episode proceeds for at most $2L$ steps. A trajectory visualisation is shown in Figure 1. Formally, we define the Graph Construction MDP (GC-MDP) as follows:

1. *State*: The state S_t is a tuple (G_t, σ_t) containing the graph $G_t = (V, E_t)$ and an *edge stub* σ_t. σ_t can be either the empty set $\varnothing$ or the singleton $\{v\}$, where $v \in V$.

2. *Action*: A_t corresponds to the selection of a node in V. Letting the degree of node v be d_v, available actions are defined as:

$$\mathcal{A}(S_t = ((V, E_t), \varnothing)) = \{v \in V \mid d_v < |V| - 1\}$$
$$\mathcal{A}(S_t = ((V, E_t), \{\sigma_t\})) = \{v \in V \mid (\sigma_t, v) \notin E_t\}$$

3. *Transitions*: The transition model is defined as $P(S_t = s' | S_{t-1} = s, A_{t-1} = a) = \delta_{S_t S'}$,

$$\text{where } S' = \begin{cases} ((V, E_{t-1} \cup (\sigma_{t-1}, a)), \varnothing), & \text{if } 2 \mid t \\ ((V, E_{t-1}), \{a\}), & \text{otherwise} \end{cases}$$

4. *Reward*: The reward R_t is defined as follows[1]:

$$R_t = \begin{cases} \mathcal{F}(G_t) - \mathcal{F}(G_0), & \text{if } t = 2L \\ 0, & \text{otherwise} \end{cases}$$

Definition of Objective Functions for Robustness. We are interested in the robustness of graphs as objective functions. Given a graph G, we let the *critical fraction* $p(G, \xi) \in [0, 1]$ be the minimum fraction of nodes that have to be removed from G in some order ξ for it to become disconnected (i.e., have more than one connected component). Connectedness is a crucial operational constraint and the higher this fraction is, the more robust the graph can be said to be. [2] The order ξ in which nodes are removed can have an impact on p, and corresponds to different scenarios: random removal is typically used to model arbitrary failures, while targeted removal is adopted as a model for attack . Formally, we consider both random permutations ξ_{random} of nodes in G, as well as permutations $\xi_{targeted}$, which are subject to the constraint that nodes must appear in the order of their degree, i.e.,

$$\forall v, u \in V. \; \xi_{targeted}(v) \leq \xi_{targeted}(u) \iff d_v \geq d_u$$

We define the objective functions $\mathcal{F}$ in the following way:

---

[1]Since $\mathcal{F}$ is very expensive to estimate, we deliberately only provide the reward at the end of the episode in order to make the training feasible computationally, to the detriment of possible credit assignment issues. Intermediate rewards based on the true objective or a related quantity represent a middle ground which we leave for future work.

[2]We note that while connectedness is required for the specific objective functions considered in this work, it is not required by either the GC-MDP formulation or the learning mechanism itself. Other robustness objectives that quantify, e.g., the size of the largest connected component are also applicable, as are fundamentally different objectives.

1. *Expected Critical Fraction to Random Removal:*

$$\mathcal{F}_{random}(G) = \mathbb{E}_{\xi_{random}}[p(G, \xi_{random})]$$

2. *Expected Critical Fraction to Targeted Removal:*

$$\mathcal{F}_{targeted}(G) = \mathbb{E}_{\xi_{targeted}}[p(G, \xi_{targeted})]$$

We use Monte Carlo (MC) sampling for estimating these quantities. For completeness, Algorithm 1 in the Supplementary Material describes how the simulations are performed. In the remainder of the paper, we use $\mathcal{F}_{random}(G)$ and $\mathcal{F}_{targeted}(G)$ to indicate their estimates obtained in this way. We highlight that evaluating an MC sample has time complexity $O(|V| \times (|V| + |E|))$: it involves checking connectedness (an $O(|V| + |E|)$ operation) after the removal of each of the $O(|V|)$ nodes. Typically, many such samples need to be used to obtain a low-variance estimate of the quantities. Coupled with the number of possible topologies, the high cost renders even shallow search methods infeasible in this domain.

3 Learning to Build Robust Graphs with Function Approximation

While the problem formulation described in Section 2 may allow us to work with a tabular RL method, the number of states quickly becomes intractable – for example, there are approximately 10^{57} labelled, connected graphs with 20 vertices [45]. Thus, we require a means of considering graph properties that are label-agnostic, permutation-invariant, and generalise across similar states and actions. Graph Neural Network architectures address these requirements. In particular, we use a graph representation based on a variant of structure2vec (S2V) [18], a GNN architecture inspired by mean field inference in graphical models. [3] Given an input graph $G = (V, E)$ where nodes $v \in V$ have feature vectors $\mathbf{x}_v$, its objective is to produce for each node v an embedding vector μ_v that captures the structure of the graph as well as interactions between neighbours. This is performed in several rounds of aggregating the features of neighbours and applying an element-wise non-linear activation function such as the rectified linear unit. For each round $k \in \{1, 2, ..., K\}$, the network simultaneously applies updates of the form:

$$\mu_v^{(k+1)} = \text{relu}\left(\theta^{(1)}\mathbf{x}_v + \theta^{(2)}\sum_{u \in \mathcal{N}(v)} \mu_u^{(k)}\right)$$

where $\mathcal{N}(v)$ is the neighbourhood of node v. We initialise embeddings with $\mu_v^{(0)} = \mathbf{0} \; \forall v \in V$, and let $\mu_v = \mu_v^{(K)}$. Once node-level embeddings are obtained, permutation-invariant embeddings for a subgraph $\mathcal{S}$ can be derived by summing

---

[3]The problem formulation does not depend on the specific GNN or RL algorithm used. While further advances developed by the community in these areas [39, 28] can be incorporated, in this paper we focus on aspects specific to the challenges of optimising the global properties of graphs.

the node-level embeddings: $\mu(\mathcal{S}) = \sum_{v_i \in \mathcal{S}} \mu_{v_i}$. The node features $\mathbf{x}_v$ are one-hot 2-dimensional vectors representing whether v is the edge stub, and their use is required to satisfy the Markovian assumption behind the MDP framework (the agent "commits" to selecting v, which is now part of its present state). Each state contains at most one edge stub.

In Q-learning [59], the agent estimates the action-value function $Q(s, a)$ introduced earlier, and derives a deterministic policy that acts greedily with respect to it. The agent interacts with the environment and updates its estimates according to:

$$Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a' \in \mathcal{A}(s')} Q(s', a') - Q(s, a)]$$

During learning, exploratory random actions are taken with probability ϵ. In the case of high-dimensional state and action spaces, approaches that use a neural network to estimate $Q(s, a)$ have been successful in a variety of domains ranging from general game-playing to continuous control [41, 37]. In particular, we use the DQN algorithm: a sample-efficient method that improves on neural fitted Q-iteration [46] by use of an experience replay buffer and an iteratively updated target network for action-value function estimation. Specifically, we use two parametrisations of the Q-function depending on whether the state S_t contains an edge stub:

$$Q(S_t = (G_t, \varnothing), A_t) = \theta^{(3)} \text{relu}\left(\theta^{(4)} \left[\mu_{A_t}, \mu(G_t)\right]\right)$$
$$Q(S_t = (G_t, \{\sigma_t\}), A_t) = \theta^{(5)} \text{relu}\left(\theta^{(6)} \left[\mu_{\sigma_t}, \mu_{A_t}, \mu(G_t)\right]\right)$$

where $[\,\cdot\,,\cdot\,]$ represents concatenation. This lets the model learn *combinations* of relevant node features (e.g., that connecting two central nodes has high Q-value). The use of GNNs has several advantages: firstly, the parameters $\Theta = \{\theta^{(i)}\}_{i=1}^6$ can be learned in a goal-directed fashion for the RL objective, allowing for flexibility in the learned representation. Secondly, the embeddings have the potential to generalise to larger graphs since they control how to combine node features of neighbours in the message passing rounds and are not restricted to graphs of a particular size. We note that the underlying S2V parameters $\theta^{(1)}, \theta^{(2)}$ are shared between the two Q-function parametrisations.

4 Experimental Setup

Learning Environment. We build a learning environment that allows for the definition of an arbitrary graph objective function $\mathcal{F}$ and provides a standardised interface for agents. Our implementation of the environment, RNet–DQN and baseline agents, and experimental suite is provided as a code repository containing Docker image blueprints that enable the reproduction of the results presented herein (up to hardware differences), including the relevant tables and figures. The instructions about how to obtain, configure, and run the code are provided in the Supplementary Material.

Baselines. We compare against the following approaches:

- *Random*: This strategy randomly selects an available action.

- *Greedy*: This strategy uses lookahead and selects the action that gives the biggest improvement in the estimated value of $\mathcal{F}$ over one edge addition.

- *Preferential*: Previous works have considered preferential additions between nodes with the two lowest degrees [9], connecting a node with the lowest degree to a random node [57] or connecting the two nodes with the lowest degree product [58], i.e., adding an edge between the vertices v, u that satisfy $\text{argmin}_{v,u} d_v \cdot d_u$. We find the latter works best in all settings tested, and refer to it as *LDP*.

- *Fiedler Vector (FV)*: The concept of Fiedler Vector was introduced by [23] and for robustness improvement by [57]. This strategy adds an edge between the vertices v, u that satisfy $\text{argmax}_{v,u} |\mathbf{y}_v - \mathbf{y}_u|$, where $\mathbf{y}$ is the Fiedler vector i.e., the eigenvector of the graph Laplacian $\mathcal{L}$ corresponding to the second smallest eigenvalue.

- *Effective Resistance (ERes)*: The concept of effective resistance was introduced by [20] and for robustness improvement as a local pairwise approximation by [58]. This strategy selects vertices v, u that satisfy $\text{argmax}_{v,u} \Omega_{v,u}$. $\Omega_{v,u}$ is defined as $(\hat{\mathcal{L}}^{-1})_{vv} + (\hat{\mathcal{L}}^{-1})_{uu} - 2(\hat{\mathcal{L}}^{-1})_{vu}$, where $\hat{\mathcal{L}}^{-1}$ is the pseudoinverse of $\mathcal{L}$.

- *Supervised Learning (SL)*: We consider a supervised learning baseline by regressing on $\mathcal{F}$ to learn an approximate $\hat{\mathcal{F}}$. We use the same S2V architecture as RNet–DQN, which we train using MSE loss instead of the Q-learning loss. To select actions for a graph G, the agent considers all graphs G' that are one edge away, selecting the one that satisfies $\text{argmax}_{G'} \hat{\mathcal{F}}$.

Neural Network Architecture. For all experiments, we use an S2V embedding vector of length 64. The neural network architecture used for RNet–DQN and SL is formed of state-action embeddings obtained using S2V followed by a multi-layer perceptron; the single output unit corresponds to the $Q(s, a)$ estimate for RNet–DQN and the predicted $\hat{\mathcal{F}}$ for SL respectively. Details of hyperparameters used for the two learning-based models are provided in the Supplementary Material.

Evaluation Protocol. We evaluate RNet–DQN and baselines both on synthetic and real-world graphs. We allow agents a number of edge additions L equivalent to a percentage τ of total possible edges. As an evaluation metric, we report the cumulative reward obtained by the agents, which quantifies the improvement in the objective function value between the final and original graphs. Specifically, in the context of the robustness metrics used, the values provided measure the difference in the expected fraction of nodes that need to be removed for the network to become disconnected. Training is performed separately for each graph family, objective function $\mathcal{F}$, and value of L. Where an agent is non-deterministic (either through intrinsic stochasticity or need for training), we repeat its evaluation (and training where applicable, starting from a different random initialisation of the network weights) to compute confidence intervals. For the learned models, we record both average and maximum performance. No hyperparameter tuning is performed due to computational budget

constraints. Details about the experimental settings are provided in the Supplementary Material.

Synthetic Graphs. We consider graphs generated through the following models:

- *Erdős–Rényi (ER)*: A graph sampled uniformly out of $\mathbf{G}^{(N,m)}$ [21]. We use $m = \frac{20}{100} * \frac{N*(N-1)}{2}$, which represents 20% of all possible edges.

- *Barabási–Albert (BA)*: A growth model where n nodes each attach preferentially to M existing nodes [4]. We use $M = 2$.

We consider graphs with $|V| = 20$, allowing agents to add a percentage of the total number of edges equal to $\tau \in \{1, 2, 5\}$, which yields $L \in \{2, 5, 10\}$. For RNet–DQN and SL, we train on a disjoint set of graphs $\mathbf{G}^{train}$. We periodically measure performance on another set $\mathbf{G}^{validate}$, storing the best model found. We use $|\mathbf{G}^{train}| = 10^4$ and $|\mathbf{G}^{validate}| = 10^2$. The performance of all agents is evaluated on a set $\mathbf{G}^{test}$ with $|\mathbf{G}^{test}| = 10^2$ generated using the ER and BA models. In order to evaluate out-of-distribution generalisation, we repeat the evaluation on graphs with up to $|V| = 100$ (only up to $|V| = 50$ for Greedy and SL due to computational cost, see next section) and scale m (for ER) and L accordingly. For non-deterministic agents, evaluation (and training, where applicable) is repeated across 50 random seeds.

Real-World Graphs. In order to evaluate our approach on real-world graphs, we consider infrastructure networks (for which robustness is a critical property) extracted from two datasets: *Euroroad* (road connections in mainland Europe and parts of Western and Central Asia [67, 33], $|V| = 1174$) and *Scigrid* (a dataset of the European power grid [40], $|V| = 1479$). We split these graphs by the country in which the nodes are located, selecting the largest connected component in case they are disconnected. We then select those with $20 \leq |V| \leq 50$, obtaining 6 infrastructure graphs for Scigrid and 8 for Euroroad. The partitioning and selection procedure yields infrastructure graphs for the following countries:

- *Euroroad*: Finland, France, Kazakhstan, Poland, Romania, Russia, Turkey, Ukraine.

- *Scigrid*: Switzerland, Czech Republic, United Kingdom, Hungary, Ireland, Sweden.

Since, in this context, the performance on individual instances matters more than generalisability, we train and evaluate on each graph separately (effectively, the sets $\mathbf{G}^{train}, \mathbf{G}^{validate}, \mathbf{G}^{test}$ all have cardinality 1 and contain the same graph). *SL* is excluded for this experiment since we consider a single network. Evaluation is repeated across 10 random seeds.

5 Results

In Table 1, we present the results of our experimental evaluation for synthetic graphs. We also display the evolution of the validation loss during training

Objective	G	L	Random	LDP	FV	ERes	Greedy	SL avg	SL best	RNet–DQN avg	RNet–DQN best
$\mathcal{F}_{random}$	BA	2	$0.018_{\pm0.001}$	0.036	0.051	0.053	0.033	$0.048_{\pm0.002}$	**0.057**	$0.051_{\pm0.001}$	**0.057**
		5	$0.049_{\pm0.002}$	0.089	0.098	0.106	0.079	$0.099_{\pm0.003}$	0.122	$0.124_{\pm0.001}$	**0.130**
		10	$0.100_{\pm0.003}$	0.158	0.176	0.180	0.141	$0.161_{\pm0.008}$	0.203	$0.211_{\pm0.001}$	**0.222**
	ER	2	$0.029_{\pm0.001}$	0.100	0.103	0.103	0.082	$0.094_{\pm0.001}$	0.100	$0.098_{\pm0.001}$	**0.104**
		5	$0.071_{\pm0.002}$	0.168	0.172	**0.175**	0.138	$0.158_{\pm0.002}$	0.168	$0.164_{\pm0.001}$	0.173
		10	$0.138_{\pm0.002}$	0.238	0.252	**0.253**	0.217	$0.221_{\pm0.005}$	0.238	$0.240_{\pm0.001}$	0.249
$\mathcal{F}_{targeted}$	BA	2	$0.010_{\pm0.001}$	0.022	0.018	0.018	0.045	$0.022_{\pm0.002}$	0.033	$0.042_{\pm0.001}$	**0.047**
		5	$0.025_{\pm0.001}$	0.091	0.037	0.077	0.077	$0.055_{\pm0.003}$	0.077	$0.108_{\pm0.001}$	**0.117**
		10	$0.054_{\pm0.003}$	0.246	0.148	0.232	0.116	$0.128_{\pm0.014}$	0.217	$0.272_{\pm0.002}$	**0.289**
	ER	2	$0.020_{\pm0.002}$	0.103	0.090	0.098	**0.149**	$0.102_{\pm0.002}$	0.118	$0.122_{\pm0.001}$	0.128
		5	$0.050_{\pm0.002}$	0.205	0.166	0.215	**0.293**	$0.182_{\pm0.008}$	0.238	$0.268_{\pm0.001}$	0.279
		10	$0.098_{\pm0.003}$	0.306	0.274	0.299	0.477	$0.269_{\pm0.016}$	0.374	$0.461_{\pm0.003}$	**0.482**

Table 1: Mean cumulative reward per episode obtained by the agents on synthetic graphs with $|V| = 20$, grouped by objective function, graph family, and number of edge additions L. Each reported value represents the improvement in the expected critical fraction between the final and initial graphs.

in Figure 2. Out-of-distribution generalisation results are shown in Figure 3. The results for real-world graphs are provided in Table 2. Additionally, Figure 4 displays examples of the original and improved topologies found by our approach.

Main Findings. We summarise our findings as follows:
RNet–DQN provides competitive performance, especially for longer action sequences. Across all settings tested, RNet–DQN performed significantly better than random. On synthetic graphs, the best model obtained the highest performance in 8 out of 12 settings tested, while the average performance is at least 89% of that of the best-performing configuration. For BA graphs, RNet–DQN obtained the best performance across all tasks tested. For ER graphs, ERes performed slightly better when considering $\mathcal{F}_{random}$; for $\mathcal{F}_{targeted}$ the greedy baseline performed better for shorter sequences. For real-world graphs, RNet–DQN obtained the best performance across all tasks.
Strategies for improving $\mathcal{F}_{random}$ are easier to learn. The performance gap between the trained model and the baselines is smaller for $\mathcal{F}_{random}$, suggesting it is less complex to learn. This is also supported by the evaluation losses monitored during training, which show performance improves and plateaus more quickly. For $\mathcal{F}_{random}$ the network with randomly initialised parameters already yields policies with satisfactory results, and training brings a small improvement. In contrast, the improvements for $\mathcal{F}_{targeted}$ are much more dramatic.
Out-of-distribution generalisation only occurs for $\mathcal{F}_{random}$. The performance on larger out-of-distribution graphs is preserved for the $\mathcal{F}_{random}$ objective, and especially for BA graphs we observe strong generalisation. The performance for $\mathcal{F}_{targeted}$ decays rapidly, obtaining worse performance than the baselines as the size increases. The poor performance of the greedy policy means the $Q(s, a)$ estimates are no longer accurate under distribution shift. There are several possible explanations, e.g., the inherent noise of estimating $\mathcal{F}_{random}$ makes the neural network more robust to outliers, or that central nodes impact message passing in larger graphs differently. We think investigating this phenomenon is a worthwhile future direction of this work, since out-of-distribution generalisation does occur for $\mathcal{F}_{random}$ and evaluating the objective functions directly is prohibitively expensive for large graphs.

Objective	Dataset	Instance	Random	LDP	FV	ERes	Greedy	RNet–DQN avg	best
$\mathcal{F}_{random}$	Euroroad	Finland	$0.080_{\pm 0.019}$	0.133	0.163	0.170	0.162	$0.161_{\pm 0.015}$	**0.189**
		France	$0.057_{\pm 0.016}$	0.149	0.181	0.163	0.151	$0.178_{\pm 0.013}$	**0.202**
		Kazakhstan	$0.107_{\pm 0.018}$	0.165	0.191	0.180	0.160	$0.179_{\pm 0.010}$	**0.203**
		Poland	$0.082_{\pm 0.033}$	0.186	0.170	0.201	0.140	$0.196_{\pm 0.014}$	**0.230**
		Romania	$0.076_{\pm 0.025}$	0.196	0.170	**0.243**	0.203	$0.207_{\pm 0.013}$	0.235
		Russia	$0.084_{\pm 0.016}$	0.135	0.224	0.157	0.187	$0.199_{\pm 0.017}$	**0.230**
		Turkey	$0.092_{\pm 0.023}$	0.191	0.198	0.198	0.191	$0.215_{\pm 0.011}$	**0.247**
		Ukraine	$0.071_{\pm 0.017}$	0.158	0.186	0.151	0.098	$0.163_{\pm 0.022}$	**0.205**
	Scigrid	Switzerland	$0.050_{\pm 0.035}$	0.191	0.160	0.174	0.182	$0.198_{\pm 0.017}$	**0.226**
		Czech Republic	$0.091_{\pm 0.020}$	0.242	0.239	0.252	0.214	$0.334_{\pm 0.020}$	**0.375**
		United Kingdom	$0.111_{\pm 0.020}$	0.263	0.273	0.290	0.224	$0.321_{\pm 0.022}$	**0.379**
		Hungary	$0.051_{\pm 0.029}$	0.176	0.179	0.175	0.117	$0.148_{\pm 0.017}$	**0.185**
		Ireland	$0.090_{\pm 0.014}$	0.208	0.211	0.213	0.177	$0.201_{\pm 0.013}$	**0.228**
		Sweden	$0.097_{\pm 0.029}$	0.187	0.213	0.195	0.197	$0.213_{\pm 0.022}$	**0.276**
$\mathcal{F}_{targeted}$	Euroroad	Finland	$0.069_{\pm 0.018}$	0.149	0.112	0.112	**0.307**	$0.273_{\pm 0.009}$	0.300
		France	$0.032_{\pm 0.019}$	0.199	0.120	0.120	0.074	$0.218_{\pm 0.006}$	**0.228**
		Kazakhstan	$0.052_{\pm 0.021}$	0.161	0.137	0.124	0.229	$0.236_{\pm 0.014}$	**0.257**
		Poland	$0.010_{\pm 0.008}$	0.101	0.114	0.084	0.108	$0.230_{\pm 0.008}$	**0.248**
		Romania	$0.029_{\pm 0.021}$	0.167	0.056	0.126	0.148	$0.238_{\pm 0.021}$	**0.270**
		Russia	$0.000_{\pm 0.000}$	0.000	0.000	0.053	0.000	$0.110_{\pm 0.036}$	**0.155**
		Turkey	$0.044_{\pm 0.021}$	0.126	0.155	0.126	0.143	$0.233_{\pm 0.018}$	**0.264**
		Ukraine	$0.031_{\pm 0.023}$	0.074	0.037	0.083	0.135	$0.164_{\pm 0.006}$	**0.178**
	Scigrid	Switzerland	$0.030_{\pm 0.024}$	0.000	0.103	0.098	0.045	$0.128_{\pm 0.006}$	**0.139**
		Czech Republic	$0.038_{\pm 0.026}$	0.116	0.116	0.116	0.163	$0.242_{\pm 0.027}$	**0.284**
		United Kingdom	$0.070_{\pm 0.047}$	0.190	0.095	0.184	0.207	$0.252_{\pm 0.027}$	**0.326**
		Hungary	$0.027_{\pm 0.028}$	**0.190**	0.000	0.129	0.143	$\mathbf{0.190_{\pm 0.000}}$	**0.190**
		Ireland	$0.047_{\pm 0.023}$	0.101	0.084	0.106	0.079	$0.259_{\pm 0.011}$	**0.288**
		Sweden	$0.061_{\pm 0.021}$	0.142	0.121	0.201	0.094	$0.232_{\pm 0.008}$	**0.261**

Table 2: Results obtained on real-world graphs, split by graph instance. Each reported value represents the improvement in the expected critical fraction between the final and initial graphs.

Performance on real-world graphs is comparatively better wrt. the baselines. This is expected since training is performed separately for each graph to be optimised.

Time Complexity. We also compare the time complexities of all approaches considered below.

- *RNet–DQN*: $O(|V| + |E|)$ operations at each step: constructing node and graph-level embeddings and, based on these embeddings, performing the forward pass in the neural network to estimate $Q(s, a)$ for all valid actions.

- *Random*: $O(1)$ for sampling, assuming the environment checks action validity.

- *Greedy*: $O(|V|^4 \times (|V| + |E|))$. The improvement in $\mathcal{F}$ is estimated for all $O(|V|^2)$ possible edges. For each edge, this involves $O(|V|)$ MC simulations. As described in Section 2, each MC simulation has complexity $O(|V| \times (|V| + |E|))$.

- *LDP*: $O(|V|^2)$: computing the product of node degrees.

- *FV, ERes*: $O(|V|^3)$, since they involve computing the eigendecomposition and the Moore-Penrose pseudoinverse of the graph Laplacian respectively (may be faster in practice).

- *SL*: $O(|V|^2 \times (|V| + |E|))$. $\hat{\mathcal{F}}$ is predicted for $O(|V|^2)$ graphs that are one edge away, then an *argmax* is taken.

Figure 2: Performance on $\mathbf{G}^{validate}$ for synthetic graphs as a function of training steps. Note the different x-axes scales for each row: more training steps are typically required for longer edge addition sequences.

It is worth noting that the analysis above does not account for the cost of training, the complexity of which is difficult to determine as it depends on many hyperparameters and the specific characteristics of the problem at hand. The approach is thus advantageous in situations in which predictions need to be made quickly, over many graphs, or the model transfers well from a cheaper training regime. We also remark that, even though the method requires an upfront cost for training, this can be seen as a constant term if the number of problem instances over which we would like to obtain predictions is large. These characteristics are shared with other emergent work that tackles combinatorial optimisation with machine learning [32, 7].

6   Related Work and Discussion

Network Resilience. Network resilience was first quantified by the average shortest path distance as a function of the number of removed nodes [2]. Analysing two scale-free communication networks, the authors found that this type of network has good robustness to random failure but is vulnerable to targeted attacks. A more extensive investigation [29] analysed the robustness of several real-world networks as well as some generated by synthetic models using a variety of attack strategies. Another area of interest is the analysis of the phase transitions of the graph in terms of connectivity under the two attack strategies [16, 17]. Optimal network topologies have also been discovered – for example, under the objective of resilience to both failures and attacks, the optimal network has a bi-modal or tri-modal degree distribution [54, 52]. There exists evidence to suggest that the topological robustness of infrastructure systems is correlated to operational robustness [49]. More broadly, the resilience of systems is highly important in structural engineering and risk management [15, 24].
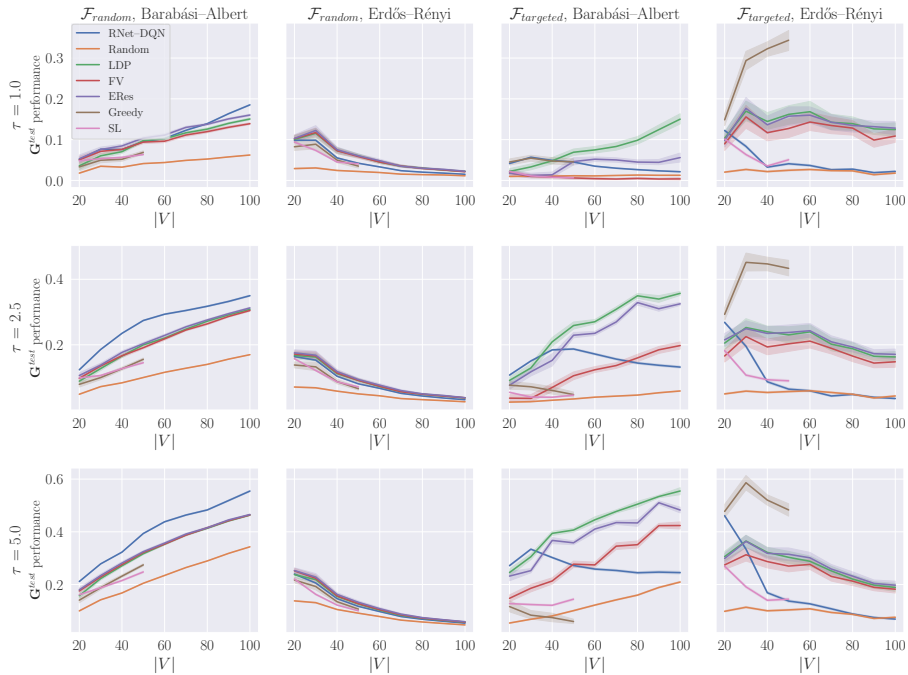
11

Figure 3: Performance on out-of-distribution synthetic graphs as a function of graph size, grouped by target problem and percentage of edge additions τ. For RNet–DQN and SL, models trained on graphs with $|V| = 20$ are used.

Graph Neural Networks and Combinatorial Optimisation. Neural network architectures able to deal not solely with Euclidean but also with manifold and graph data have been developed in recent years [12], and applied to a variety of problems where their capacity for representing structured and relational information can be exploited [5]. A sub-category of such approaches are Message Passing Neural Networks (MPNN) [26], often referred to as Graph Neural Networks (GNN) instead. Significant progress has been achieved in machine learning for combinatorial optimisation problems [7], such as Minimum Vertex Cover and the Travelling Salesman Problem by framing them as a supervised learning [56] or RL [6] task. Combining GNNs with RL algorithms has yielded models capable of solving several graph optimisation problems with the same architecture while generalising to graphs an order of magnitude larger than those used during training [32]. The problem as formulated in this paper is a combinatorial optimisation problem, related to problems in network design that arise in operations research [1] – albeit with a different objective.

Graph Generation. Similarities exist between the current work and the area of graph generative modelling, which tries to learn a generative model of graphs [35, 65, 36] in a computationally efficient way given some existing examples. This generation, however, is not necessarily conditioned by an objective that captures a global property to be optimised. Other lines of work target constructing, *ab initio*, graphs that have desirable properties: examples include neural architecture search [66, 38] and molecular graph generation [30, 64, 11].

Road Network, France

$\mathcal{F}_{random}(G) = 0.077$
$\mathcal{F}_{targeted}(G) = 0.037$

$\mathcal{F}_{random}(G) = 0.232$

$\mathcal{F}_{targeted}(G) = 0.256$

Road Network, Russia

$\mathcal{F}_{random}(G) = 0.067$
$\mathcal{F}_{targeted}(G) = 0.033$

$\mathcal{F}_{random}(G) = 0.242$

$\mathcal{F}_{targeted}(G) = 0.186$

Power Grid, Czech Republic

$\mathcal{F}_{random}(G) = 0.080$
$\mathcal{F}_{targeted}(G) = 0.023$

$\mathcal{F}_{random}(G) = 0.425$

$\mathcal{F}_{targeted}(G) = 0.278$

Figure 4: Several examples of the solutions found by RNet–DQN on real-world graphs. Each row of the illustration shows the original network on the left, while the central and right panels show the network optimised for resilience to random and targeted removals, respectively. Objective function values are shown underneath. The solutions for $\mathcal{F}_{random}$ typically assign more connections to a few central nodes, notably discovering the hub pattern in the third example. For $\mathcal{F}_{targeted}$ the added edges are spread around the network, reducing the impact of attacks. However, the algorithm might discover more complex patterns that are not directly interpretable, as shown in the solutions for the first example network.

The concurrent work GraphOpt [53] tackles the *inverse problem*: given a graph, the goal is to learn the underlying objective function that leads to its generation. This is achieved with maximum entropy inverse RL and GNNs.

13

Relationship to RL–S2V. Our work builds on RL–S2V, a method that was applied for the construction of adversarial examples against graph classifiers [19]. However, it is worth noting that there are a series of key differences with respect to that approach. First, RL–S2V is not designed to address the problem of constructing robust graphs or, more generally, learning to construct graphs according to a given goal. Secondly, there are two key algorithmic differences to RL–S2V with respect to the GC-MDP formulation: the reward function used, which in this case quantifies a global property of the graph itself, as well as the definition of the action spaces and the transition model, which account for excluding already-existing edges (RL–S2V ignores this, leading to some of the edge budget being wasted). Since the values of structural properties we consider are increasing in the number of edges (the complete graph has robustness 1), RNet–DQN generally yields strictly better performance results.

7 Conclusion and Outlook

In this work, we have addressed the problem of improving a graph structure given the goal of maximising the value of a global objective function. We have framed it for the first time as a decision-making problem and we have formalised it as the Graph Construction MDP (GC-MDP). Our approach, named *RNet–DQN*, uses Reinforcement Learning and Graph Neural Networks as key components for generalisation. As a case study, we have considered the problem of improving graph robustness to random and targeted removals of nodes. Our experimental evaluation on synthetic and real-world graphs shows that, in certain situations, this approach can deliver performance superior to existing methods, both in terms of the solutions found (i.e., the resulting robustness of the graphs) and time complexity of model evaluation. Further, we have shown the ability to transfer to out-of-sample graphs, as well as the potential to transfer to out-of-distribution graphs larger than those used during training.

Extensions. The proposed approach can be applied to other problems based on different definitions of robustness or considering fundamentally different objective functions such as efficiency [34], path diversity [27], and assortativity [42], which are of interest in various biological, communication, and social networks. Since our formulation and algorithm are objective-agnostic, we expect they are applicable out-of-the-box for other objectives, even those for which no strong baselines are currently known. As such, this approach may be a useful tool for the discovery of new graph improvement algorithms for objectives that can be evaluated programatically, either in closed form or via simulations. Potential limitations might be related to the complexity of the objective functions and the related computational demands. We also view the interpretability of the approach, which is less straightforward than those based on known mathematical concepts such as the Fiedler vector, to be an important research direction. Since there is an active interest in the interpretability of both GNNs and RL [63, 55], we consider that there is scope for developing techniques that are tailor-made for explaining policies learned by RL on graphs.

Operationalisation. Beyond considering other objectives, in order to operationalise the proposed algorithm, it is possible to integrate a variety of re-

finements, which can include capturing heterogeneous edge costs (e.g., different capacities per link in a communication network), extending the action space to support heterogenous edge types (e.g., addition of different types of edges with specific characteristics), and integrating domain-specific link constraints (e.g., planarity). For critical scenarios, it is also possible to verify that the resulting solutions satisfy some given formal properties and constraints [25]. Furthermore, considering multi-criteria objective functions [47] is important for cases where properties of the solutions must be balanced [31]. Various choices exist for representing this trade-off and should be captured on a case-by-case basis depending on the application: for example, a linear combination may be sufficient in certain situations, while others are characterised by economies of scale. We also remark that our formulation captures operational scenarios in which the cost of constructing a link is significantly greater than the cost of its maintenance (e.g., as with road networks). Our method can also be adapted for situations in which operational cost is instead greater by formulating a rewiring operation that keeps the number of links constant.

Broader Applications.   The approach described in this work can be used to improve the properties of a variety of human-made infrastructure systems, such as communication networks, transportation networks, and power grids. We also envisage potential applications in biological (e.g., hypothesis testing for understanding the characteristics of brain networks [13]), ecological (e.g., design of more resilient ecosystems [61]) and social networks (e.g., design of organisational structures [62]). As far as biological networks are concerned, the brain is hypothesised to optimise a trade-off between efficiency and wiring cost [13]. Our method could be used in order to test different hypotheses related to the resulting structure of brain networks over evolutionary times and also during their development. Indeed, the evolution of such networks in time has been captured (for example, Sulston et al. [51] mapped the development of the C. elegans connectome). This can be achieved by applying the optimisation procedure for different objective functions and comparing the obtained networks to the "ground truth". With respect to the potential application in ecosystem management, the graph formalisation can be used to model interactions between species in a given environment. As such, our method has potential applications to study, in simulation, the impact of introducing or removing species from an ecosystem so as to achieve a desired outcome. Specifically, in the context of robustness, we can consider optimising the resilience of an ecosystem to intrinsic or extrinsic shocks, a task of fundamental importance [3]. The dynamics of interactions between species may be modelled in simulation using well known models of e.g., predator-prey mechanics [8]. With respect to social networks, for example, our method can be applied to derive optimal communication strategies and related team structures so as to optimise a given objective for an organisation. Finally, there are also potential applications for networks of artificial agents (i.e., robots). There is a significant body of work in the robotics literature that treats the problem of maintaining robust communication in a network of agents working together to complete a task in an environment that contains obstacles or adversaries. For instance, [50] uses properties of the graph Laplacian (namely, the Fiedler vector and its associated eigenvalue) to ensure the underlying communication network remains robust. Since we have empirically

15

shown superior performance to using the Fiedler vector, our approach could also lead to gains in this deployment scenario.

References

[1] Ravindra K. Ahuja, Thomas L. Magnanti, James B. Orlin, and M. R. Reddy. Chapter 1 Applications of network optimization. In *Handbooks in Operations Research and Management Science*, volume 7 of *Network Models*, pages 1–83. Elsevier, 1995.

[2] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, 2000.

[3] Stefano Allesina, Antonio Bodini, and Mercedes Pascual. Functional links and robustness in food webs. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1524):1701–1709, 2009.

[4] Albert-László Barabási and Réka Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999.

[5] Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Álvaro Sánchez-González, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv:1806.01261*, 2018.

[6] Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural Combinatorial Optimization with Reinforcement Learning. *arXiv:1611.09940*, 2016.

[7] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine Learning for Combinatorial Optimization: a Methodological Tour d'Horizon. *European Journal of Operational Research*, 290:405–421, 2021.

[8] Alan A Berryman. The origins and evolution of predator-prey theory. *Ecology*, 73(5):1530–1535, 1992.

[9] Alina Beygelzimer, Geoffrey Grinstein, Ralph Linsker, and Irina Rish. Improving Network Robustness by Edge Modification. *Physica A*, 357:593–612, 2005.

[10] Monica Bianchini, Marco Gori, and Franco Scarselli. Inside PageRank. *ACM Trans. Internet Technol.*, 5(1):92–128, February 2005.

[11] John Bradshaw, Brooks Paige, Matt J. Kusner, Marwin H. S. Segler, and José Miguel Hernández-Lobato. A Model to Search for Synthesizable Molecules. In *NeurIPS*, 2019.

[12] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric Deep Learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

[13] Ed Bullmore and Olaf Sporns. The economy of brain network organization. *Nature Reviews Neuroscience*, 13(5):336–349, 2012.

[14] Hale Cetinay, Karel Devriendt, and Piet Van Mieghem. Nodal vulnerability to targeted attacks in power grids. *Applied Network Science*, 3(1):34, 2018.

[15] Gian Paolo Cimellaro, Andrei M. Reinhorn, and Michel Bruneau. Framework for analytical quantification of disaster resilience. *Engineering Structures*, 32:3639–3649, 2010.

[16] Reuven Cohen, Keren Erez, Daniel ben Avraham, and Shlomo Havlin. Resilience of the Internet to Random Breakdowns. *Physical Review Letters*, 85(21):4626–4628, 2000.

[17] Reuven Cohen, Keren Erez, Daniel ben Avraham, and Shlomo Havlin. Breakdown of the Internet under Intentional Attack. *Physical Review Letters*, 86(16):3682–3685, 2001.

[18] Hanjun Dai, Bo Dai, and Le Song. Discriminative embeddings of latent variable models for structured data. In *ICML*, 2016.

[19] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. Adversarial attack on graph structured data. In *ICML*, 2018.

[20] W. Ellens, F.M. Spieksma, P. Van Mieghem, A. Jamakovic, and R.E. Kooij. Effective graph resistance. *Linear Algebra and its Applications*, 435(10):2491–2506, 2011.

[21] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

[22] Ernesto Estrada and Naomichi Hatano. Communicability in complex networks. *Physical Review E*, 77(3):036111, 2008.

[23] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

[24] Alexander A. Ganin, Emanuele Massaro, Alexander Gutfraind, Nicolas Steen, Jeffrey M. Keisler, Alexander Kott, Rami Mangoubi, and Igor Linkov. Operational resilience: Concepts, design and analysis. *Scientific Reports*, 6(1):1–12, 2016.

[25] Javier García and Fernando Fernández. A Comprehensive Survey on Safe Reinforcement Learning. *Journal of Machine Learning Research*, 16(42):44, 2015.

[26] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural Message Passing for Quantum Chemistry. In *ICML*, 2017.

[27] Nikola Gvozdiev, Stefano Vissicchio, Brad Karp, and Mark Handley. On low-latency-capable topologies, and their impact on the design of intra-domain routing. In *SIGCOMM*, 2018.

[28] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining Improvements in Deep Reinforcement Learning. In *AAAI*, 2018.

[29] Petter Holme, Beom Jun Kim, Chang No Yoon, and Seung Kee Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5), 2002.

[30] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction Tree Variational Autoencoder for Molecular Graph Generation. In *ICML*, 2018.

[31] Ralph L. Keeney, Howard Raiffa, and Richard F. Meyer. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge University Press, 1993.

[32] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization algorithms over graphs. In *NeurIPS*, 2017.

[33] Jérôme Kunegis. KONECT: the Koblenz network collection. In *WWW Companion*, 2013.

[34] Vito Latora and Massimo Marchiori. Efficient Behavior of Small-World Networks. *Physical Review Letters*, 87(19):198701, 2001.

[35] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning Deep Generative Models of Graphs. In *ICML*, 2018.

[36] Renjie Liao, Yujia Li, Yang Song, Shenlong Wang, Charlie Nash, William L. Hamilton, David Duvenaud, Raquel Urtasun, and Richard S. Zemel. Efficient Graph Generation with Graph Recurrent Attention Networks. In *NeurIPS*, 2019.

[37] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, et al. Continuous control with deep reinforcement learning. In *ICLR*, 2016.

[38] Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. Hierarchical representations for efficient architecture search. In *ICLR*, 2018.

[39] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably Powerful Graph Networks. In *NeurIPS*, 2019.

[40] Wided Medjroubi, Ulf Philipp Müller, Malte Scharf, Carsten Matke, and David Kleinhans. Open Data in Power Grid Modelling: New Approaches Towards Transparent Grid Models. *Energy Reports*, 3:14–21, 2017.

[41] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[42] M. E. J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(20), 2002.

[43] M. E. J. Newman. The Structure and Function of Complex Networks. *SIAM Review*, 45(2), 2003.

[44] M. E. J. Newman. *Networks*. Oxford University Press, 2018.

[45] OEIS Foundation. The on-line encyclopedia of integer sequences, 2020.

[46] Martin Riedmiller. Neural Fitted Q Iteration – First Experiences with a Data Efficient Neural Reinforcement Learning Method. In *ECML*, 2005.

[47] D. M. Roijers, P. Vamplew, S. Whiteson, and R. Dazeley. A Survey of Multi-Objective Sequential Decision-Making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.

[48] Christian M. Schneider, André A. Moreira, Joao S. Andrade, Shlomo Havlin, and Hans J. Herrmann. Mitigation of malicious attacks on networks. *PNAS*, 108(10):3838–3841, 2011.

[49] Ricard V. Solé, Martí Rosas-Casals, Bernat Corominas-Murtra, and Sergi Valverde. Robustness of the European power grids under intentional attack. *Physical Review E*, 77(2):026102, 2008.

[50] Ethan Stump, Ali Jadbabaie, and Vijay Kumar. Connectivity management in mobile robot teams. In *2008 IEEE International Conference on Robotics and Automation*, pages 1525–1530. IEEE, 2008.

[51] John E Sulston, Einhard Schierenberg, John G White, and J Nichol Thomson. The embryonic cell lineage of the nematode caenorhabditis elegans. *Developmental Biology*, 100(1):64–119, 1983.

[52] Toshi Tanizawa, Gerald Paul, Reuven Cohen, Shlomo Havlin, and H. Eugene Stanley. Optimization of network robustness to waves of targeted and random attacks. *Physical Review E*, 71(4), 2005.

[53] Rakshit Trivedi, Jiachen Yang, and Hongyuan Zha. GraphOpt: Learning Optimization Models of Graph Formation. In *ICML*, 2020.

[54] André X. C. N. Valente, Abhijit Sarkar, and Howard A. Stone. Two-Peak and Three-Peak Optimal Complex Networks. *Physical Review Letters*, 92(11), 2004.

[55] Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, and Swarat Chaudhuri. Programmatically interpretable reinforcement learning. In *ICML*, 2018.

[56] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer Networks. In *NeurIPS*, 2015.

[57] Huijuan Wang and Piet Van Mieghem. Algebraic connectivity optimization via link addition. In *Proceedings of the Third International Conference on Bio-Inspired Models of Network Information and Computing Systems (Bionetics)*, 2008.

[58] Xiangrong Wang, Evangelos Pournaras, Robert E. Kooij, and Piet Van Mieghem. Improving robustness of complex networks via the effective graph resistance. *The European Physical Journal B*, 87(9):221, 2014.

[59] Christopher J. C. H. Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3-4):279–292, 1992.

[60] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440, 1998.

[61] Walter E. Westman. Measuring the inertia and resilience of ecosystems. *BioScience*, 28(11):705–710, 1978.

[62] John Wreathall. Properties of resilient organizations: an initial view. In *Resilience Engineering*, pages 275–285. CRC Press, 2017.

[63] Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. Gnnexplainer: Generating explanations for graph neural networks. In *NeurIPS*, 2019.

[64] Jiaxuan You, Bowen Liu, Rex Ying, Vijay Pande, and Jure Leskovec. Graph Convolutional Policy Network for Goal-Directed Molecular Graph Generation. In *NeurIPS*, 2018.

[65] Jiaxuan You, Rex Ying, Xiang Ren, William L. Hamilton, and Jure Leskovec. GraphRNN: Generating Realistic Graphs with Deep Autoregressive Models. In *ICML*, 2018.

[66] Barret Zoph and Quoc V. Le. Neural Architecture Search with Reinforcement Learning. In *ICLR*, 2017.

[67] L. Šubelj and M. Bajec. Robust network community detection using balanced propagation. *The European Physical Journal B*, 81(3):353–362, 2011.