

Exploratory Combinatorial Optimization with Reinforcement Learning

Thomas D. Barrett,¹ William R. Clements,² Jakob N. Foerster,³ A. I. Lvovsky^{1,4}

¹University of Oxford, Oxford, UK

²Uncharted Technologies, Paris, France

³Facebook AI Research

⁴Russian Quantum Center, Moscow, Russia

thomas.barrett@physics.ox.ac.uk, william.clements@unchartech.com, jnf@fb.com, alex.lvovsky@physics.ox.ac.uk

Abstract

Many real-world problems can be reduced to combinatorial optimization on a graph, where the subset or ordering of vertices that maximize some objective function must be found. With such tasks often NP-hard and analytically intractable, reinforcement learning (RL) has shown promise as a framework with which efficient heuristic methods to tackle these problems can be learned. Previous works construct the solution subset incrementally, adding one element at a time, however, the irreversible nature of this approach prevents the agent from revising its earlier decisions, which may be necessary given the complexity of the optimization task. We instead propose that the agent should seek to continuously improve the solution by learning to *explore at test time*. Our approach of exploratory combinatorial optimization (ECO-DQN) is, in principle, applicable to any combinatorial problem that can be defined on a graph. Experimentally, we show our method to produce state-of-the-art RL performance on the Maximum Cut problem. Moreover, because ECO-DQN can start from any arbitrary configuration, it can be combined with other search methods to further improve performance, which we demonstrate using a simple random search.

1 Introduction

NP-hard combinatorial problems – such as Travelling Salesman (Papadimitriou 1977), Minimum Vertex Cover (Dinur and Safra 2005) and Maximum Cut (Goemans and Williamson 1995) – are canonical challenges in computer science. With applications across numerous practical settings, ranging from fundamental science to industry, efficient methods for approaching combinatorial optimization are of great interest. However, as no known algorithms are able to solve NP-hard problems in polynomial time, exact methods rapidly become intractable for all but the simplest of tasks. Approximation algorithms guarantee a worst-case solution quality, but sufficiently strong bounds may not exist and, even if they do, these algorithms can have limited scalability (Williamson and Shmoys 2011). Instead, heuristics are often deployed that, despite offering no theoretical guarantees, are chosen for high performance.

There are numerous heuristic methods, ranging from search-based (Benlic and Hao 2013; Banks, Vincent, and

Anyakoha 2008) to physical systems that utilise both quantum and classical effects (Johnson et al. 2011; Yamamoto et al. 2017) and their simulated counterparts (Kirkpatrick, Gelatt, and Vecchi 1983; Clements et al. 2017; Tiunov, Ulanov, and Lvovsky 2019). However, the effectiveness of general algorithms is dependent on the problem being considered, and high levels of performance often require extensive tailoring and domain-specific knowledge. Machine learning offers a route to addressing these challenges, which led to the demonstration of a meta-algorithm, S2V-DQN (Khalil et al. 2017), that utilises reinforcement learning (RL) and a deep graph network to automatically learn good heuristics for various combinatorial problems.

A solution to a combinatorial problem defined on a graph consists of a subset of vertices that satisfies the desired optimality criteria. S2V-DQN, and the related works discussed shortly, incrementally construct solutions one element at a time – reducing the problem to predicting the value of adding any vertex not currently in the solution to this subset. However, due to the inherent complexity of many combinatorial problems, learning a policy that directly produces a single, optimal solution is often impractical, as evidenced by the sub-optimal performance of such approaches.

Instead, we propose that a natural reformulation is for the agent to explore the solution space at test time, rather than producing only a single “best-guess”. Concretely, this means the agent can add or remove vertices from the solution subset and is tasked with searching for ever-improving solutions at test time. In this work we present ECO-DQN (Exploratory Combinatorial Optimization DQN), a framework combining RL and deep graph networks to realise this approach.

Our experimental work considers the Maximum Cut (Max-Cut) problem as it is a fundamental combinatorial challenge – in fact over half of the 21 NP-complete problems enumerated in Karp’s seminal work (Karp 1972) can be reduced to Max-Cut – with numerous real-world applications. The framework we propose can, however, be readily applied to any graph-based combinatorial problem where solutions correspond to a subset of vertices and the goal is to optimize some objective function.

By comparing ECO-DQN to S2V-DQN as a baseline, we demonstrate that our approach improves on the state-of-the-art for applying RL to the Max-Cut problem. With suitable ablations we further show that this performance gap is de-

ECO-DQN is available at <https://github.com/tomdbar/eco-dqn>, including code to reproduce all experiments reported in this work.

pendent on both allowing the agent to reverse its earlier decisions and providing suitable information and rewards to exploit this freedom. Moreover, ECO-DQN can be initialised in any state (i.e. will look to improve on any proposed solution) and, as a consequence, has the flexibility to be either deployed independently or combined with other search heuristics. For example, we achieve significant performance improvements by simply taking the best solution found across multiple randomly initialised episodes.

Our agents are trained to find solutions for any arbitrary graph from a given distribution, however we also show that ECO-DQN generalises well to graphs from unseen distributions. We obtain very strong performance on known benchmarks of up to 2000 vertices, even when the agent is trained on graphs an order of magnitude smaller and with a different structure.

Related Work

A formative demonstration of neural networks for combinatorial optimization (CO) was the application of Hopfield networks to the Travelling Salesman Problem (TSP) by Hopfield and Tank (1985). They mapped N -city problems to $N \times N$ graphs (networks) with each vertex, v_{ij} , a binary variable denoting whether city i is the j -th to be visited, and the edges connecting vertices proportional to the distance between cities. Although the results of Hopfield and Tank were contested by Wilson and Pawley (1988), there followed a period active research into neural networks for CO that lasted for over a decade (Smith 1999).

RL techniques were applied to CO for the first time by Bello *et al.* (2016). They used policy gradients to train pointer networks (Vinyals, Fortunato, and Jaitly 2015), a recurrent architecture that produces a softmax attention mechanism (a “pointer”) to select a member of the input sequence as an output. However, this architecture did not reflect the structure of problems defined over a graph, which Khalil *et al.* (2017) addressed with S2V-DQN, a general RL-based framework for CO that uses a combined graph embedding network and deep Q-network. Mittal *et al.* (2019) developed these ideas further by modifying the training process: first training an embedding graph convolution network (GCN), and then training a Q-network to predict the vertex (action) values. This is orthogonal to our proposal which considers the framework itself, rather than the training procedure, and, in principle, appears to be compatible with ECO-DQN.

Another current direction is applying graph networks for CO in combination with a tree search. Li *et al.* (2018) combined a GCN with a guided tree-search in a supervised setting, i.e. requiring large numbers of pre-solved instances for training. Very recently, Abe *et al.* (2019) trained a GCN using Monte-Carlo tree search as a policy improvement operator, in a similar manner to AlphaGo Zero (Silver *et al.* 2017), however, this work does not consider the Max-Cut problem.

2 Background

Max-Cut Problem

The Max-Cut problem is to find a subset of vertices on a graph that maximises the total number of edges connect-

ing vertices within this subset to vertices not in this subset (the cut value). In this work we consider the more general weighted version of this problem, where each edge in the graph is assigned a weight and the objective is to maximise the total value of cut edges. Formally, for a graph, $G(V, W)$, with vertices V connected by edges W , the Max-Cut problem is to find the subset of vertices $S \subset V$ that maximises $C(S, G) = \sum_{i \in S, j \in V \setminus S} w_{ij}$ where $w_{ij} \in W$ is the weight of the edge connecting vertices i and j .

This is not simply a mathematical challenge as many real world applications can be reduced to the Max-Cut problem, including protein folding (Perdomo-Ortiz *et al.* 2012), investment portfolio optimization (Elsokkary *et al.* 2017; Venturelli and Kondratyev 2018) (specifically using the Markowitz (1952) formulation), and finding the ground state of the Ising Hamiltonian in physics (Barahona 1982).

Q-learning

As is standard for RL, we consider the optimization task as a Markov decision process (MDP) defined by the 5-tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$. Here, \mathcal{S} denotes the set of states, \mathcal{A} is the set of actions, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition function, $\mathcal{R} : \mathcal{S} \rightarrow \mathbb{R}$ the reward function and $\gamma \in [0, 1]$ is the discount factor. A policy, $\pi : \mathcal{S} \rightarrow [0, 1]$, maps a state to a probability distribution over actions. The Q-value of a given state-action pair, $(s \in \mathcal{S}, a \in \mathcal{A})$, is then given by the discounted sum of immediate and future rewards

$$Q^\pi(s, a) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t) \mid s_0 = s, a_0 = a, \pi \right], \quad (1)$$

where s_0 and a_0 correspond to the initial state and action taken, with future actions chosen according to the policy, π .

A deep Q-network (Mnih *et al.* 2015) (DQN) provides a function $Q(s, a; \theta)$, where θ parameterises the network, which is trained to approximate $Q^*(s, a) \equiv \max_{\pi} Q^\pi(s, a)$, the Q-values of each state-action pair when following the optimal policy. Once trained, an approximation of the optimal policy can be obtained simply by acting greedily with respect to the predicted Q-values, $\pi(s; \theta) = \operatorname{argmax}_{a'} Q(s, a'; \theta)$.

Message Passing Neural Networks

Our choice of deep Q-network is a message passing neural network (MPNN) (Gilmer *et al.* 2017). This is a general framework of which many common graph networks are specific implementations. The basic idea is to represent each vertex in the graph, $v \in V$, with some n -dimensional embedding, μ_v^k , where k labels the current iteration (network layer). These embeddings are initialised according to

$$\mu_v^0 = I(x_v), \quad (2)$$

where I is some initialisation function and $x_v \in \mathbb{R}^m$ is the input vector of observations. During the message-passing phase, the embeddings are repeatedly updated with information from neighbouring vertices according to

$$m_v^{k+1} = M_k \left(\mu_v^k, \{ \mu_u^k \}_{u \in N(v)}, \{ w_{uv} \}_{u \in N(v)} \right), \quad (3)$$

$$\mu_v^{k+1} = U_k \left(\mu_v^k, m_v^{k+1} \right), \quad (4)$$

where M_k and U_k are message and update functions, respectively, with $N(v)$ is the set of vertices directly connected to v . After K rounds of message passing, a prediction – a set of values that carry useful information about the network – is produced by some readout function, R . In our case this prediction is the set of Q-values of the actions corresponding to “flipping” each vertex, i.e. adding or removing it from the solution subset S ,

$$\{Q_v\}_{v \in V} = R\left(\{\mu_u^K\}_{u \in V}\right). \quad (5)$$

3 Exploiting Exploration

One straightforward application of Q-learning to CO over a graph is to attempt to directly learn the utility of adding any given vertex to the solution subset. This formalism, which is followed by S2V-DQN and related works, incrementally constructs a solution by adding one vertex at a time to the subset, until no further improvement can be made. However, the complexity of NP-hard combinatorial problems means it is challenging to learn a single function approximation of $Q^*(s, a)$ that generalises across the vast number of possible graphs. Therefore, as vertices can only be added to the solution set, policies derived from the learnt Q-values, such as a typical greedy policy, will likely be sub-optimal.

In this work we present an alternative approach where the agent is trained to explore the solution space at test time, seeking ever-improving states. As such, the Q-value of either adding or removing a vertex from the solution is continually re-evaluated in the context of the episode’s history. Additionally, as all actions can be reversed, the challenge of predicting the true value of a vertex “flip” does not necessarily result in sub-optimal performance. The fundamental change distinguishing our approach, ECO-DQN, from previous works can then be summarised as follows: *instead of learning to construct a single good solution, learn to explore for improving solutions.*

However, simply allowing for revisiting the previously flipped vertices does not automatically improve performance. The agent is not immediately able to make more informed decisions, nor can it reach previously unobtainable solutions. Instead, further modifications are required to leverage this freedom for improved performance, which we discuss here.

Reward Shaping. The objective of our exploring agent is to find the best solution (highest cut-value) at any point within an episode. Formally, the reward at state $s_t \in \mathcal{S}$ is given by $\mathcal{R}(s_t) = \max(C(s_t) - C(s^*), 0)/|V|$, where $s^* \in \mathcal{S}$ is the state corresponding to the highest cut value previously seen within the episode, $C(s^*)$ (note that we implicitly assume the graph, G , and solution subset, S , to be included in the state). As continued exploration is desired, even after a good solution is found, there is no punishment if a chosen action reduces the cut-value. Note also that the reward is normalised by the total number of vertices, $|V|$, to mitigate the impact of different reward scales across different graph sizes. We use a discount factor of $\gamma = 0.95$ to ensure the agent actively pursues rewards within a finite time horizon.

Our environment only provides a reward when a new best solution is found and, as a result, after an initial period of exploration, these extrinsic rewards can be relatively sparse, or even absent, for the remainder of the episode. We therefore also provide a small intermediate reward of $1/|V|$ whenever the agent reaches a locally optimal state (one where no action will immediately increase the cut value) previously unseen within the episode. In addition to mitigating the effect of sparse extrinsic rewards, these intrinsic rewards also shape the exploratory behaviour at test time. As there are far more states than could be visited within our finite episodes, the vast majority of which are significantly sub-optimal, it is useful to focus on a subset of states known to include the global optimum. As local optima in combinatorial problems are typically close to each other, the agent learns to “hop” between nearby local optima, thereby performing a in-depth local search of the most promising subspace of the state space (see figure 2b).

Observations A Q-value for flipping each vertex is calculated using seven observations derived from the current state ($x_v \in \mathbb{R}^7$). These observations are:


1. Vertex state, i.e. if v is currently in the solution set, S .
2. Immediate cut change if vertex state is changed.
3. Steps since the vertex state was last changed.
4. Difference of current cut-value from the best observed.
5. Distance of current solution set from the best observed.
6. Number of available actions that immediately increase the cut-value.
7. Steps remaining in the episode.

Observations (1-3) are local, which is to say they can be different for each vertex considered, whereas (4-7) are global, describing the overall state of the graph and the context of the episode. The general purposes of each of the observations are: (1-2) provide useful information for determining the value of selecting an action, (3) provides a simple history to prevent short looping trajectories, (4-5) ensure the rewards calculated with respect to the best observed cut-value are Markovian, (6) allows the agent to predict when it will receive the intrinsic rewards previously discussed and (7) accounts for the finite episode duration.

Experiments

Experimental details. In this work we train and test the agent on both Erdős-Rényi (Erdős and Rényi 1960) and Barabasi-Albert (Albert and Barabási 2002) graphs with edges $w_{ij} \in \{0, \pm 1\}$, which we refer to as ER and BA graphs, respectively. Training is performed on randomly generated graphs from either distribution, with each episode considering a freshly generated instance. The performance over training (i.e. all learning curves) is evaluated as the mean of a fixed set of 50 held-out graphs from the same distribution. Once trained, the agents are tested on a separate set of 100 held-out validation graphs from a given distribution.

During training and testing, every action taken demarks a timestep, t . For agents that are allowed to take the same action multiple times (i.e. ECO-DQN and selected ablations),



(a) Learning curves: ER graphs with $|V|=40$

Agent	$ V =20$	$ V =40$	$ V =60$	$ V =100$	$ V =200$
ER graphs					
ECO-DQN	$0.97^{+0.03}_{-0.03}$	$1.00^{+0.00}_{-0.00}$	$0.99^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.01}$
S2V-DQN	$0.97^{+0.03}_{-0.03}$	$0.98^{+0.01}_{-0.02}$	$0.98^{+0.01}_{-0.02}$	$0.92^{+0.02}_{-0.02}$	$0.95^{+0.02}_{-0.02}$
MCA-irrev	$0.89^{+0.06}_{-0.11}$	$0.89^{+0.04}_{-0.05}$	$0.87^{+0.05}_{-0.05}$	$0.87^{+0.03}_{-0.04}$	$0.86^{+0.03}_{-0.03}$
BA graphs					
ECO-DQN	$0.99^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.01}$	$0.98^{+0.00}_{-0.02}$	$0.97^{+0.02}_{-0.03}$	$0.93^{+0.02}_{-0.03}$
S2V-DQN	$0.97^{+0.01}_{-0.03}$	$0.96^{+0.03}_{-0.04}$	$0.94^{+0.02}_{-0.04}$	$0.95^{+0.02}_{-0.03}$	$0.94^{+0.02}_{-0.02}$
MCA-irrev	$0.92^{+0.05}_{-0.08}$	$0.89^{+0.05}_{-0.06}$	$0.88^{+0.04}_{-0.05}$	$0.87^{+0.03}_{-0.04}$	$0.87^{+0.03}_{-0.03}$

(b) Single episode performance: ER and BA graphs

Figure 1: (a) Learning curves, averaged over 5 seeds, of agents training on 40-vertex ER graphs. ECO-DQN is compared to S2V-DQN as a baseline, with the differences individually ablated as described in the text. (b) Approximation ratios of ECO-DQN, S2V-DQN and the MCA-irrev heuristics for ER and BA graphs with different numbers of vertices, $|V|$. We report the mean approximation ratio of each agent over the 100 validation graphs, along with the distance to the upper and lower quartiles as a guide to how varied the performance is across different graph instances.

which for convenience we will refer to as *reversible* agents, the episode lengths are set to twice the number of vertices in the graph, $t = 1, 2, \dots, 2|V|$. Each episode for such agents is initialised with a random subset of vertices in the solution set. By contrast, agents that can only add vertices to the solution set (*irreversible* agents, i.e. S2V-DQN and selected ablations) are initialised with an empty solution subset. These agents keep selecting actions greedily even if no positive Q-values are available until $t = |V|$, to account for possible incorrect predictions of the Q-values. In both cases the best solution obtained at any timestep within the episode is taken as the final result. To facilitate direct comparison, ECO-DQN and S2V-DQN are implemented with the same MPNN architecture, with details provided in the Appendix.

Benchmarking details We compare the performance of ECO-DQN to a leading RL-based heuristic, S2V-DQN. To interpret the performance gap, we also consider the following ablations, which together fully account for the differences between our approach and the baseline (ECO-DQN \equiv S2V-DQN+RevAct+ObsTun+IntRew).

- *Reversible Actions* (RevAct): Whether the agent is allowed to flip a vertex more than once. For irreversible agents we follow S2V-DQN and use $\gamma=1$.
- *Observation Tuning* (ObsTun): Observations (2-7) from the list above that allow the agent to exploit having reversible actions. Also, in the absence of ObsTun, the rewards used are simply the (normalised) immediate change in cut value, $\mathcal{R}(s_t) = (C(s_t) - C(s_{t-1}))/|V|$, which is necessary as without observations (4-5) the ECO-DQN reward structure is non-Markovian.
- *Intermediate Rewards* (IntRew): Whether the agent is provided with the small intermediate rewards for reaching new locally optimal solutions.

As an additional benchmark we also implement the Max-CutApprox (MCA) algorithm. This is a greedy algorithm, choosing the action (vertex) that provides the greatest immediate increase in cut value until no further improvements can be made. We consider two modifications of MCA. The

standard application, which we denote MCA-irrev, is irreversible and begins with an empty solution set. The alternative algorithm, MCA-rev, starts with a random solution set and allows reversible actions.

We use the approximation ratio of each approach as a metric of solution quality. This is defined as $\alpha = C(s^*)/C(s_{\text{opt}})$, where $C(s_{\text{opt}})$ is the cut-value associated with the true optimum solution. Exact methods are intractable for many of the graph sizes we use, therefore we apply a battery of optimization approaches to each graph and take the best solution found by any of them as the “optimum” solution. Specifically, in addition to ECO-DQN, S2V-DQN and the MCA algorithms, we use CPLEX, an industry standard integer programming solver, and a pair of recently developed simulated annealing heuristics by Tiunov *et al.* (2019) and Leleu *et al.* (2019). Details of these implementations and a comparison of their efficacy can be found in the Supplemental Material.

Performance benchmarking. Figure 1a shows learning curves of agents trained on ER graphs of size $|V| = 40$, where it can be seen that ECO-DQN reaches a significantly higher average cut than S2V-DQN. We see that removing either reversible actions (RevAct) or the additional observations (ObsTun) reduces the performance below that of S2V-DQN, underlining our previous assertion that obtaining state-of-the-art performance requires not only that the agent be allowed to reverse its previous actions, but also that it be suitably informed and rewarded to do so effectively. The intermediate rewards (IntRew) can be seen to speed up and stabilise training. They also result in a small performance improvement, however this effect becomes clearer when considering how the agents generalise to larger graphs (see figures 3a and 3b).

Figure 1b shows the performance of agents trained and tested on both ER and BA graphs of sizes ranging from 20 to 200 vertices. We see that ECO-DQN has superior performance across most considered graph sizes and structures.

Intra-episode behaviour. We now consider how this strong performance is achieved by examining the intra-

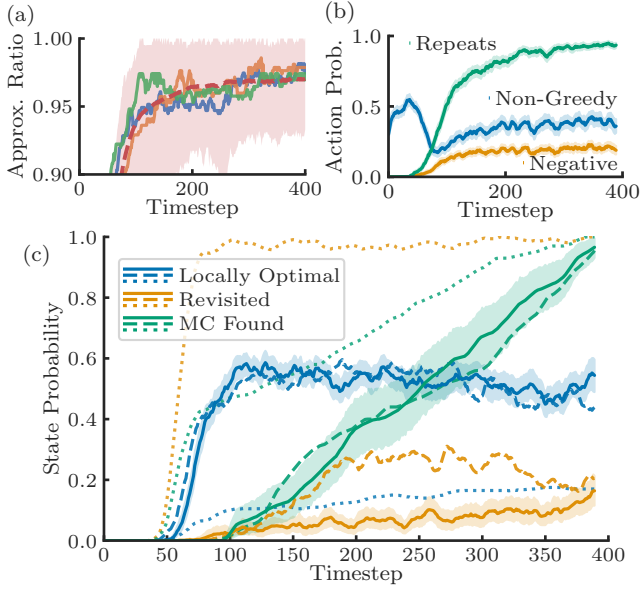


Figure 2: The intra-episode behaviour of an agent trained and tested on 200 vertex ER graphs with 400 actions per episodes. (a) The trajectories on the three graphs in the validation set, with the mean and range of all trajectories plotted as the dashed line and shaded region, respectively. (b) The probability that an action taken by the agent has already been taken within the episode (Repeats), does not provide the greatest immediate reward (Non-Greedy) or reduces the cut value (Negative). (c) Plots the probability that each state visited is locally optimal (Locally Optimal) or has already been visited within the episode (Revisited). Also shown is the probability at each timestep that the best solution that will be found within the episode has already been seen (MC found). The behaviour of the agent is shown at three points during training: when performance is equivalent to that of either MCA-irrev (dotted lines) or S2V-DQN (dashed) and when the agent is fully trained (solid). Both (b) and (c) use a 10-step moving average over all 100 graphs (trajectories) in the validation set, however, for clarity, the shaded errors are only shown for the fully trained agent in (c).

episode behaviour of an agent trained and tested on ER graphs with 200 vertices (figure 2). The larger graph size is chosen as it provides greater scope for the agent to exhibit non-trivial behaviour.

Figure 2a highlights the trajectories taken by the trained agent on three random graphs. We observe that the trajectories, whilst trending towards higher cut values, have short term fluctuations. This illustrates that the agent has learnt how to search for improving solutions even if it requires short-term sacrificing of the cut value. Further analysis of the agent’s behaviour is presented in figures 2b and 2c which show the action preferences and the types of states visited, respectively, over the course of an optimization episode.

From figure 2b, we see that the fully trained agent regularly chooses actions that do not correspond to the greatest immediate increase in the cut-value (Non-Greedy), or even

that decrease the cut value (Negative). Moreover, the agent also moves the same vertex in or out of the solution set multiple times within an episode (Repeats), which suggests the agent has learnt to explore multiple possible solutions that may be different from those obtained initially.

This is further emphasised in figure 2c where we see that, after an initial period of exploration, the agent searches through the solution space, repeatedly moving in and out of locally optimal (Locally Optimal) solutions whilst minimising the probability that it revisits states (Revisited). By comparing the agent’s behaviour at three points during training (fully trained and when performance level is equivalent to either MCA-irrev or S2V-DQN), we see that this behaviour is learnt. Weaker agents from earlier in training revisit the same states far more often, yet find fewer locally optimal states. To show that this exploratory behaviour improves the agent’s performance, we also plot the probability that the agent has already found the best solution it will see in the episode (MC found). We see this probability (averaged over 100 graphs) grow monotonically, implying that the agent keeps finding ever better solutions while exploring. This suggests that longer episode lengths could lead to even better performance. Indeed, for this agent, simply increasing the number of timesteps in an episode from $2|V|=400$ to $4|V|$ is seen to increase the average approximation ratio from $0.98^{+0.01}_{-0.01}$ to $0.99^{+0.01}_{-0.01}$. As in figure 1b, the range quoted for these approximation ratios corresponds to the upper and lower quartiles of the performance across all 100 validation graphs.

4 Leveraging Variance

Changing the initial subset of vertices selected to be in the solution set can result in very different trajectories over the course of an episode. An immediate result of this stochasticity is that performance can be further improved by running multiple episodes with a distribution of initialisations, and selecting the best result from across this set, as we show in this section.

Experiments

We optimize every graph using 50 randomly initialised episodes. At the same time, we make the task more challenging by testing on graphs that are larger, or that have a different structure, from those on which the agent was trained. This ability to generalise to unseen challenges is important for the real-world applicability of RL agents to combinatorial problems where the distribution of optimization tasks may be unknown or even change over time.

Generalisation to unseen graph types. Figures 3a and 3b show the generalisation of agents trained on 40 vertices to systems with up to 500 vertices for ER and BA graphs, respectively. (Generalisation data for agents trained on graphs of sizes ranging from $|V| = 20$ to $|V| = 200$ can be found in the Appendix.) ECO-DQN is compared to multiple benchmarks, with details provided in the caption, however there are three important observations to emphasise. Firstly, reversible agents outperform the irreversible benchmarks on all tests, with the performance gap widening with increasing

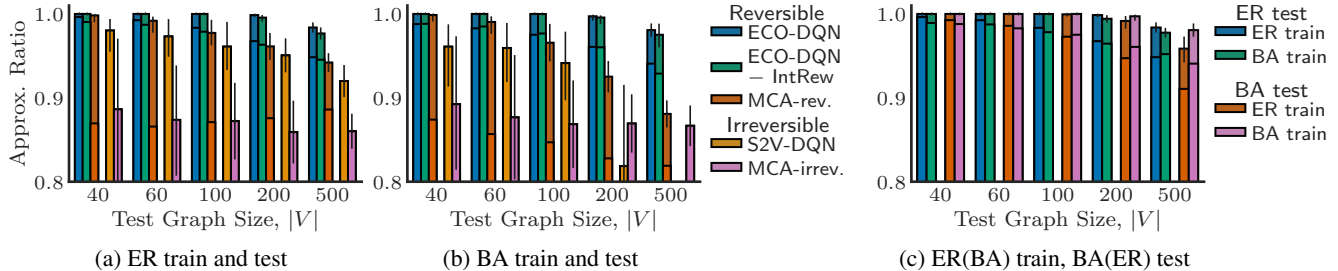


Figure 3: Generalisation of agents to unseen graph sizes and structures. (a-b) The performance of agents trained on ER and BA graphs of size $|V|=40$ tested on graphs of up to $|V|=500$ of the same type. ECO-DQN is shown with and without providing the intermediate rewards (IntRew) for finding locally optimal solutions during training, and is compared to the MCA-rev algorithm. Reversible agents are applied with 50 randomly initialised episodes to each of the 100 validation graphs for each type and size. The first marking on each bar is the average across every episode (the expected ‘single-try’ performance), with the upper limit extending to the average performance across different graphs. The vertical bars denote the 68 % confidence interval of this upper limit. Irreversible approaches are initialised with empty solution sets, and so only use 1 episode per graph. Note that S2V-DQN applied to BA graphs with $|V|=500$ is not visible on these axes but has a value of $0.49^{+0.12}_{-0.11}$. (c) A comparison of how agents trained on only one of either ER or BA graphs with $|V|=40$, perform on larger graphs from both distributions.

graph size. This is particularly noticeable for BA graphs, for which the degrees of each vertex in the graph tend to be distributed over a greater range than for ER graphs, where S2V-DQN fails to generalise in any meaningful way to $|V| \geq 200$.

Secondly, for the reversible agents it is clear that using multiple randomly initialised episodes provides a significant advantage. As ECO-DQN provides near-optimal solutions on small graphs within a single episode, it is only on larger graphs that this becomes relevant. However, it is noteworthy that even the simple MCA-rev algorithm, with only a relatively modest budget of 50 random initialisations, outperforms a highly trained irreversible heuristic (S2V-DQN). This further emphasises how stochasticity – which here is provided by the random episode initialisations and ensures many regions of the solution space are considered – is a powerful attribute when combined with local optimization.

Finally, we again observe the effect of small intermediate rewards (IntRew) for finding locally optimal solutions during training upon the final performance. For small graphs the agent performs near-optimally with or without this intrinsic motivation, however the difference becomes noticeable when generalising to larger graphs at test time.

In figure 3c we observe that ECO-DQN performs well across a range of graph structures, even if they were not represented during training, which is a highly desirable characteristic for practical CO. We train the agent on ER graphs with $|V|=40$ and then test it on BA graphs of up to $|V|=500$, and visa versa. The performance is marginally better when testing on graphs from the same distribution as the training data, however this difference is negligible for $|V| \leq 100$. Furthermore, in every case deploying ECO-DQN with 50 randomly initialised episodes outperforms all other benchmarks (S2V-DQN and MCA), even when only ECO-DQN is trained on different graph types to the test data.

Generalization to real-world datasets. Finally, we test ECO-DQN on publicly available datasets. The first is the “Physics” dataset consisting of ten graphs, with $|V|=125$,

which correspond to Ising spin-glass models of physical systems. This dataset consists of regular graphs with exactly 6 connections per vertex and $w_{ij} \in \{0, \pm 1\}$. The second is the GSet, a benchmark collection of large graphs that have been well investigated (Benlic and Hao 2013). We separately consider the first ten graphs, G1-G10 which have $|V|=800$, and the first ten larger graphs, G22-G32, which have $|V|=2000$. For G1-G10 we utilise 50 randomly initialised episodes per graph, however for G22-G32 we use only a single episode per graph, due to the increased computational cost. We apply agents trained on ER graphs with $|V|=200$. The results are summarised in table 1, where ECO-DQN is seen to significantly outperform other approaches, even when restricted to use only a single episode per graph.

Dataset ($ V $)	ECO-DQN	S2V-DQN	MCA-(rev, irrev)
Physics (125)	1.000	0.928	0.879, 0.855
G1-10 (800)	0.996	0.950	0.947, 0.913
G22-32 (2000)	0.971	0.919	0.883, 0.893

Table 1: Average performance on known benchmarks.

Despite the structure of graphs in the “Physics” dataset being distinct from the ER graphs, also with $w_{ij} \in \{0, \pm 1\}$, on which the agent is trained, every instance is solved with, on average, a 37.6 % chance of a given episode finding the optimal solution. Moreover, 90.4 % of the optimal solutions found are unique, demonstrating that, in conjunction with random initialisations, the agent is capable of finding many different optimal trajectories. The structure of the GSet is also distinct from that of the training data, with the first five instances in each tested set have only positive edges.

5 Summary and Outlook

This work introduces ECO-DQN, a new state-of-the-art RL-based algorithm for the Max-Cut problem that generalises well to unseen graph sizes and structures. We show that

treating CO as an ongoing exploratory exercise in surpassing the best observed solution is a powerful approach to this NP-hard problem. In principle, our approach is applicable to any combinatorial problem defined on a graph.

ECO-DQN can initialise a search from any valid state, opening the door to combining it with other search heuristics. This feature, as we demonstrate, leads to further improved performance even with this “heuristic” being as simple as a set of randomly initialised episodes. However, one could consider combining ECO-DQN with more sophisticated episode initialisation policies. Alternatively, ECO-DQN could also be initialised with solutions found by other optimization methods to further strengthen them. Also, we train our agents with highly discounted future rewards ($\gamma = 0.95$), and although this is found to provide strong performance, the relatively short-term reward horizon likely limits exploration to only local regions of the solution space. As such, it would be interesting to investigate longer reward horizons, particularly when training on larger graphs. A more substantial avenue to explore would be to use a recurrent architecture where a useful representation of the episode history is learned, as opposed to the hand-crafted representation that we describe in section 3.

Whilst these are paths towards further developing exploration-based CO, we believe that the strong performance already demonstrated would allow our approach to be applied in numerous practical settings. This is especially true for settings where many graphs of similar structure need to be optimized, such as protein folding (Perdomo-Ortiz et al. 2012) and portfolio optimization (Elsokkary et al. 2017; Venturelli and Kondratyev 2018).

Acknowledgements

The authors would like to thank D. Chermoshentsev and A. Boev for their expertise in applying simulated annealing heuristics and CPLEX to our validation graphs. A.L.’s research is partially supported by Russian Science Foundation (19-71-10092).

Appendix

Graph Instances. The training, testing and validation graphs were generated with the NetworkX Python package (Hagberg, Schult, and Swart 2008). For ER graphs, a connection probability of 0.15 is used. The BA graphs have an average degree of 4. To produce our target graphs we then randomly set all non-zero edges to ± 1 .

The established benchmarks from table 1 are publicly available and will also be included with the released code for this work.

MPNN architecture. The initial embedding for each vertex, v , is given by

$$\mu_v^0 = \text{relu}(\theta_1 x_v) \quad (6)$$

where $x_v \in \mathbb{R}^m$ is the input vector of observations and $\theta_1 \in \mathbb{R}^{m \times n}$. We also learn embeddings describing the connections to each vertex v ,

$$\xi_v = \text{relu}\left(\theta_3 \left[\frac{1}{|N(v)|} \sum_{u \in N(v)} \text{relu}(\theta_2[w_{uv}, x_u]), |N(v)|\right]\right), \quad (7)$$

where $\theta_2 \in \mathbb{R}^{m+1 \times n-1}$, $\theta_3 \in \mathbb{R}^{n \times n}$ and square bracket denote concatenation. The embeddings at each vertex are then updated according to

$$m_v^{k+1} = \text{relu}\left(\theta_{4,k} \left[\frac{1}{|N(v)|} \sum_{u \in N(v)} w_{uv} \mu_u^k, \xi_v\right]\right), \quad (8)$$

$$\mu_v^{k+1} = \text{relu}\left(\theta_{5,k} [\mu_v^k, m_v^{k+1}]\right), \quad (9)$$

where $\{\theta_{4,k}, \theta_{5,k}\} \in \mathbb{R}^{2n \times n}$. After K rounds of message passing, the Q-value for a vertex is read out using the final embeddings across the entire graph,

$$Q_v = \theta_7 \left[\text{relu}\left(\theta_6 \frac{1}{|V|} \sum_{u \in V} \mu_u^K\right), \mu_v^K\right], \quad (10)$$

with $\theta_6 \in \mathbb{R}^{n \times n}$ and $\theta_7 \in \mathbb{R}^{2n}$.

In this work we use $n=64$ dimensional embedding vectors, and have $K=3$ rounds of message passing. However, many different MPNN implementations can be used with good success. In general, what is important is that the network be capable of capturing relevant information about the local neighbourhood of a vertex.

Training Details. All agents are trained with a minibatch sizes of 64 and 32 actions per step of gradient descent. The learning rate is 10^{-4} and the exploration rate is linearly decreased from $\varepsilon=1$ to $\varepsilon=0.05$ over the first $\sim 10\%$ of training. We use the same MPNN for both S2V-DQN and ECO-DQN. We verify that this network properly represents S2V-DQN by reproducing its performance on the ‘Physics’ dataset at the level reported in the original work by Khalil *et al.* (2017). During the training of an irreversible agent’s Q-network, the predictions of the Q-values produced by the target network are clipped to be strictly non-negative. This clipping is also used by Khalil *et al.* and is empirically observed to improve and stabilise training.

Test \rightarrow Train \downarrow	$ V =20$	$ V =40$	$ V =60$	$ V =100$	$ V =200$	$ V =500$	$ V =20$	$ V =40$	$ V =60$	$ V =100$	$ V =200$	$ V =500$
	ER graphs						BA graphs					
$ V =20$	$0.99^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$0.98^{+0.01}_{-0.01}$	$0.95^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$0.99^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.01}$
$ V =40$	—	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$0.98^{+0.01}_{-0.01}$	—	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$0.98^{+0.01}_{-0.01}$
$ V =60$	—	—	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$0.99^{+0.01}_{-0.01}$	—	—	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$0.99^{+0.01}_{-0.01}$
$ V =100$	—	—	—	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	—	—	—	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$0.98^{+0.01}_{-0.01}$
$ V =200$	—	—	—	—	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	—	—	—	—	$0.99^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.01}$

Table 2: Generalisation performance of ECO-DQN, using 50 randomly initialised episodes per graph.

Full Generalisation Data. ECO-DQN’s generalisation performance on ER and BA graphs is shown in table 2.

References

- Abe, K.; Xu, Z.; Sato, I.; and Sugiyama, M. 2019. Solving NP-Hard Problems on Graphs by Reinforcement Learning without Domain Knowledge. *arXiv:1905.11623*.
- Albert, R., and Barabási, A.-L. 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics* 74(1):47.
- Banks, A.; Vincent, J.; and Anyakoha, C. 2008. A review of particle swarm optimization. Part II: hybridisation, combinatorial, multicriteria and constrained optimization, and indicative applications. *Natural Computing* 7(1):109–124.
- Barahona, F. 1982. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General* 15(10):3241.
- Bello, I.; Pham, H.; Le, Q. V.; Norouzi, M.; and Bengio, S. 2016. Neural Combinatorial Optimization with Reinforcement Learning. *arXiv:1611.09940*.
- Benlic, U., and Hao, J.-K. 2013. Breakout Local Search for the Max-Cut problem. *Engineering Applications of Artificial Intelligence* 26(3):1162 – 1173.
- Clements, W. R.; Renema, J. J.; Wen, Y. H.; Chrzanowski, H. M.; Kolthammer, W. S.; and Walmsley, I. A. 2017. Gaussian Optical Ising Machines. *Phys. Rev. A* 96:043850.
- Dinur, I., and Safra, S. 2005. On the Hardness of Approximating Minimum Vertex Cover. *Annals of Mathematics* 439–485.
- Elsokkary, N.; Khan, F. S.; La Torre, D.; Humble, T. S.; and Gottlieb, J. 2017. Financial Portfolio Management using D-Wave Quantum Optimizer: The Case of Abu Dhabi Securities Exchange. Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States).
- Erdős, P., and Rényi, A. 1960. On the Evolution of Random Graphs. *Publ. Math. Inst. Hung. Acad. Sci* 5(1):17–60.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural Message Passing for Quantum Chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1263–1272.
- Goemans, M. X., and Williamson, D. P. 1995. Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems using Semidefinite Programming. *Journal of the ACM (JACM)* 42(6):1115–1145.
- Hagberg, A. A.; Schult, D. A.; and Swart, P. J. 2008. Exploring Network Structure, Dynamics, and Function using NetworkX. In *Proceedings of the 7th Python in Science Conference*.
- Hopfield, J. J., and Tank, D. W. 1985. “Neural” Computation of Decisions in Optimization Problems. *Biological Cybernetics* 52(3):141–152.
- Johnson, M. W.; Amin, M. H.; Gildert, S.; Lanting, T.; Hamze, F.; Dickson, N.; Harris, R.; Berkley, A. J.; Johansson, J.; Bunyk, P.; et al. 2011. Quantum Annealing with Manufactured Spins. *Nature* 473(7346):194.
- Karp, R. M. 1972. Reducibility Among Combinatorial Problems. In *Complexity of Computer Computations*. Springer. 85–103.
- Khalil, E.; Dai, H.; Zhang, Y.; Dilkina, B.; and Song, L. 2017. Learning Combinatorial Optimization Algorithms over Graphs. In *Advances in Neural Information Processing Systems*, 6348–6358.
- Kirkpatrick, S.; Gelatt, C. D.; and Vecchi, M. P. 1983. Optimization by Simulated Annealing. *Science* 220(4598):671–680.
- Kochenberger, G. A., and Glover, F. 2006. A unified framework for modeling and solving combinatorial optimization problems: A tutorial. In *Multiscale Optimization Methods and Applications*. Springer. 101–124.
- Leleu, T.; Yamamoto, Y.; McMahon, P. L.; and Aihara, K. 2019. Destabilization of local minima in analog spin systems by correction of amplitude heterogeneity. *Phys. Rev. Lett.* 122:040607.
- Li, Z.; Chen, Q.; and Koltun, V. 2018. Combinatorial Optimization with Graph Convolutional Networks and Guided Tree Search. In *Advances in Neural Information Processing Systems*, 539–548.
- Markowitz, H. 1952. Portfolio Selection. *The Journal of Finance* 7(1):77–91.
- Mittal, A.; Dhawan, A.; Medya, S.; Ranu, S.; and Singh, A. 2019. Learning Heuristics over Large Graphs via Deep Reinforcement Learning. *arXiv:1903.03332*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Papadimitriou, C. H. 1977. The Euclidean travelling salesman problem is NP-complete. *Theoretical Computer Science* 4(3):237–244.
- Perdomo-Ortiz, A.; Dickson, N.; Drew-Brook, M.; Rose, G.; and Aspuru-Guzik, A. 2012. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific Reports* 2:571.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of Go without Human Knowledge. *Nature* 550(7676):354.
- Smith, K. A. 1999. Neural Networks for Combinatorial Optimization: A Review of More Than a Decade of Research. *INFORMS Journal on Computing* 11(1):15–34.
- Tiunov, E. S.; Ulanov, A. E.; and Lvovsky, A. 2019. Annealing by simulating the coherent Ising machine. *Optics Express* 27(7):10288–10295.
- Venturelli, D., and Kondratyev, A. 2018. Reverse Quantum Annealing Approach to Portfolio Optimization Problems. *Quantum Machine Intelligence* 1–14.
- Vinyals, O.; Fortunato, M.; and Jaitly, N. 2015. Pointer Networks. In *Advances in Neural Information Processing Systems*, 2692–2700.
- Wigley, P. B.; Everitt, P. J.; van den Hengel, A.; Bastian, J. W.; Sooriyabandara, M. A.; McDonald, G. D.; Hardman, K. S.; Quinlivan, C. D.; Manju, P.; Kuhn, C. C.; et al. 2016. Fast machine-learning online optimization of ultra-cold-atom experiments. *Scientific reports* 6:25890.
- Williamson, D. P., and Shmoys, D. B. 2011. *The Design of Approximation Algorithms*. Cambridge University Press.
- Wilson, G., and Pawley, G. 1988. On the stability of the Traveling Salesman Problem algorithm of Hopfield and Tank. *Biological Cybernetics* 58(1):63–70.
- Yamamoto, Y.; Aihara, K.; Leleu, T.; Kawarabayashi, K.-i.; Kako, S.; Fejer, M.; Inoue, K.; and Takesue, H. 2017. Coherent Ising machines—optical neural networks operating at the quantum limit. *npj Quantum Information* 3(1):49.

Supplemental Material

We apply six different optimization methods to the 100 validation graphs of each structure (ER or BA graphs) and size ($|V| \in \{20, 40, 60, 100, 200, 500\}$). The highest cut value across the board is then chosen as the reference point that we refer to as the “optimum value”. Table 3 compares the performance of these methods on our validation sets. The number of times each method reaches these “optimum” solutions is shown in table 4. Here, we give an overview of each method and summarise their efficacy.

ECO-DQN. The framework introduced and discussed in detail in the main text. We train distinct agents on every graph structure and size, up to $|V| = 200$, and then test them on graphs of the same or larger sizes. For each individual agent-graph pair, we run 50 randomly initialised optimization episodes. Therefore, graphs with $|V| = 20$ are subject to only 50 optimization attempts, whereas graphs with $|V| = 500$ are optimised with 300 episodes using 5 distinct agents. The performance of each agent is summarised in table 2 of the main text.

S2V-DQN. An RL framework for graph-based combinatorial problems introduced by Khalil *et al.* (2017). Details of our implementation can, again, be found in the main text. S2V-DQN agents are trained and tested equivalently to the ECO-DQN agents. However, as S2V-DQN is deterministic at test time, only a single optimization episode is used for every agent-graph pair.

MCA. The final optimization method introduced in the main text is MaxCutApprox (MCA). This is a simple greedy algorithm that can be applied either in the reversible or irreversible setting. We refer to these as MCA-rev and MCA-irrev, respectively. For every optimization episode of ECO-DQN or S2V-DQN, a corresponding MCA-rev or MCA-irrev episode is also undertaken. We take the best solution found in any episode by either of these greedy algorithms as the MCA solution.

We see from tables 3 and 4 that the greedy MCA algorithms find optimal solutions on nearly all graphs of size up to $|V| = 60$, but performance rapidly deteriorates thereafter. As the number of possible solution configurations (states) grows exponentially with the number of vertices, this simply reflects how it quickly becomes intractable to sufficiently cover the state-space in our finite number of episodes.

$ V \rightarrow$ Approach \downarrow	20	40	60	100	200	500
ER graphs						
ECO-DQN	0.99	1.00	1.00	1.00	1.00	1.00
S2V-DQN	0.97	0.99	0.99	0.98	0.96	0.95
MCA	1.00	1.00	1.00	0.99	0.98	0.96
CPLEX	1.00	1.00	1.00	0.87	0.46	0.16
SimCIM	1.00	1.00	1.00	1.00	0.99	0.99
Leleu <i>et al.</i>	1.00	1.00	1.00	1.00	1.00	1.00
BA graphs						
ECO-DQN	1.00	1.00	1.00	1.00	1.00	0.99
S2V-DQN	0.97	0.98	0.98	0.97	0.96	0.92
MCA	1.00	1.00	1.00	0.99	0.95	0.90
CPLEX	1.00	1.00	1.00	1.00	0.83	0.17
SimCIM	1.00	1.00	1.00	0.99	0.99	0.97
Leleu <i>et al.</i>	1.00	1.00	1.00	1.00	0.94	1.00

Table 3: The approximation ratios, averaged across 100 graphs for each graph structure and size, of the different optimization methods.

CPLEX. A commercial optimizer. We first transform the Max-Cut problem into a QUBO (Quadratic Unconstrained Binary Optimization) task (Kochenberger and Glover 2006). Strictly, for a graph, $G(V, W)$, with vertices V connected by edges W , this task is to minimize the Hamiltonian

$$H = - \sum_{i,j} w_{ij} (x_i - x_j)^2, \quad (11)$$

where $x_k \in \{\pm 1\}$ labels whether vertex $k \in V$ is in the solution subset, $S \subset V$. This Hamiltonian is then solved using mixed integer programming by the CPLEX branch-and-bound routine.

For each graph, we take the best solution found within 10 min as the final answer. Within this time budget we find the exact solution on all graphs with up to 60 vertices. Only some of the 100-vertex graphs are optimally solved, with performance significantly dropping for the 200 and 500 vertex graphs due to the unfeasibly large solution space.

SimCIM. A simulated annealing heuristic proposed by Tiunov *et al.* (2019) that models the classical dynamics within a coherent Ising machine (CIM) (Yamamoto *et al.* 2017). This approach relaxes the binary vertex values ($x_k \in \{\pm 1\}$), associated with labelling a vertex as either in or out of the solution subset, to analog values ($-1 \leq x_k \leq 1$). Each vertex is initialised to 0, and then subjected to evolution according to a set of stochastic differential equations that describe the operation of the CIM. In the process of the evolution, the system eventually settles with all vertices in near-binary states. Details of both CIM and SimCIM beyond the high-level description given here can be found in the referenced works. The hyperparameters of SimCIM were optimised using a differential evolution approach by M-LOOP (Wigley *et al.* 2016) over 50 runs.

Leleu *et al.* Another recently developed simulated annealing heuristic that relaxes the binary vertex labels to analog values. A key feature of this approach is the modification of the time-dependent interaction strengths in such a way as to destabilise locally optimal solutions. Details can be found in the work of Leleu *et al.* (2019). As with SimCIM, the hyperparameters are adjusted by M-LOOP (Wigley *et al.* 2016) over 50 runs.

$ V \rightarrow$ Approach \downarrow	20	40	60	100	200	500
ER graphs						
ECO-DQN	99	100	100	100	100	50
S2V-DQN	76	67	55	18	0	0
MCA	100	100	93	49	3	0
CPLEX	100	100	100	0	0	0
SimCIM	100	92	92	71	13	0
Leleu <i>et al.</i>	100	100	100	98	100	92
BA graphs						
ECO-DQN	100	100	100	100	95	12
S2V-DQN	74	60	39	18	1	0
MCA	100	99	85	24	0	0
CPLEX	100	100	100	97	0	0
SimCIM	100	87	94	60	12	0
Leleu <i>et al.</i>	100	97	99	100	96	97

Table 4: The relative contributions of the different optimization methods to the “optimum” solutions. Shown is the number of graphs (out of 100) for which each approach finds the best, or equal best, solution.