

Lawrence Berkeley National Laboratory

LBL Publications

Title

Graph Partitioning and Sparse Matrix Ordering using Reinforcement Learning and Graph Neural Networks

Permalink

<https://escholarship.org/uc/item/08z43090>

Authors

Gatti, Alice
Hu, Zhixiong
Smidt, Tess
[et al.](#)

Publication Date

2021-04-07

Peer reviewed

Graph Partitioning and Sparse Matrix Ordering using Reinforcement Learning and Graph Neural Networks

Alice Gatti* Zhixiong Hu[†] Tess Smidt* Esmond G. Ng* Pieter Ghysels*

June 30, 2021

Abstract

We present a novel method for graph partitioning, based on reinforcement learning and graph convolutional neural networks. Our approach is to recursively partition coarser representations of a given graph. The neural network is implemented using SAGE graph convolution layers, and trained using an advantage actor critic (A2C) agent. We present two variants, one for finding an edge separator that minimizes the normalized cut or quotient cut, and one that finds a small vertex separator. The vertex separators are then used to construct a nested dissection ordering to permute a sparse matrix so that its triangular factorization will incur less fill-in. The partitioning quality is compared with partitions obtained using METIS and SCOTCH, and the nested dissection ordering is evaluated in the sparse solver SuperLU. Our results show that the proposed method achieves similar partitioning quality as METIS and SCOTCH. Furthermore, the method generalizes across different classes of graphs, and works well on a variety of graphs from the SuiteSparse sparse matrix collection.

1 Introduction

The problem of partitioning a graph into approximately equal sized subgraphs while minimizing the number of cut edges is an NP-complete problem [19]. Practical graph partitioning algorithms for large scale problems are based on heuristics which give approximate solutions. In general it is not known how far these approximations are from the optimal solution. Moreover, many of the heuristic algorithms are inherently sequential in nature and do not exploit current high-performance computing hardware efficiently. Heuristic algorithms exhibit very irregular memory access patterns – leading to low memory bandwidth usage because of wasted cache lines and branch prediction misses – and perform mostly integer manipulations. Deep learning is an expressive and flexible algorithmic framework that can run efficiently on modern hardware, making it a powerful tool for tackling the problem of graph partitioning from a new perspective.

In recent years there has been a growing interest in using machine learning techniques to find approximate solutions of NP problems, like the travelling salesman problem, knapsack problem, vertex cover etc. Many of the problems in this class are naturally formulated as combinatorial optimization problems over graphs, and powerful learning tools to handle this type of data are graph neural networks [9, 22, 7], which are an extension of usual deep neural networks to non-Euclidean data such as graphs. One of the key ingredients of graph neural networks is graph convolution, that generalizes convolution over grids to graph data structures. This generalization is often called message passing and it allows to aggregate local information on the neighbors and propagate it on the graph. This scheme has been integrated in many different convolutional layers such as GCN [31], SAGE [21] and GAT [47], in which message passing is combined with the attention mechanism. Several machine learning methods including graph neural networks have been proposed to approximate the solutions of the travelling salesman problem [51, 8, 42], knapsack problem [2], vertex cover and maximum cut [29]. These works propose supervised and unsupervised methods together with reinforcement learning. There are also works focused on the graph partitioning problem. [25] combines a probabilistic method with graph neural networks, but it has not been tested on big graphs. [38] presents an unsupervised deep learning

*Computational Research Division, Lawrence Berkeley National Laboratory

[†]University of California, Santa Cruz




Figure 2: Multilevel approach for graph partitioning. On the top of the picture the graph is coarsened recursively, until the coarsest graph as few vertices. The coarsening is made by matching vertices included in the same gray ellipse. Then, on the bottom part of the figure, the coarsest graph is partitioned and it is interpolated back by one level, where the partition is further refined. The interpolation/refinement continues up to the initial graph. The vertices included in the red or blue ellipses have been interpolated from a red or blue node respectively.

Although the KL and FM algorithms are widely used in practice, they have limited parallelism and have irregular memory access patterns and lots of integer manipulations, making them inefficient on modern hardware. Spectral graph methods on the other hand, are based on eigenvalue solvers and can take advantage of the huge computational power of modern CPUs, including SIMD units, and GPUs. Still, spectral methods can be expensive, as the number of iterations for the eigensolver, for instance Lanczos [34], Rayleigh quotient iteration [46] or LOBPCG [32], can be large. Popular graph partitioning libraries are METIS [27] and SCOTCH [39], and their parallel versions ParMETIS [26] and PT-SCOTCH [11], as well as Zoltan(2) [14].

Graphs arise frequently in scientific computing. For example, a mesh used in the discretization of partial differential equations can be considered a graph. Sparse matrices are another example; their sparsity structures can be represented by graphs as well. Partitioning of graphs is an important pre-processing step in high-performance parallel computing, with the goal of distributing the computational work evenly over the compute nodes while minimizing communication in, for instance, Krylov iterative solvers and preconditioners like algebraic multigrid (AMG) [16], block Jacobi and domain decomposition [44]. For instance the widely used scientific computing codes PETSc [6], Trilinos [24] and the MFEM finite element library [4] all use METIS for graph partitioning.

The remainder of the paper is organized as follows. Section 2 briefly introduces reinforcement learning, in particular the distributed advantage actor critic (DA2C) training algorithm. In Section 3 we present an algorithm to compute a minimal edge separator using deep reinforcement learning within a multilevel framework. Section 4 shows a variation of this algorithm presented to compute a vertex separator instead of an edge separator. The vertex separator algorithm is used in Section 5 to construct a nested dissection sparse matrix ordering. We conclude the paper with a summary and outlook in Section 6. All the codes are made available at the GitHub page <https://github.com/alga-hopf/drl-graph-partitioning>.

2 Advantage Actor Critic

Reinforcement learning [45] is an extremely flexible framework for training an agent, interacting with an environment, to maximize its cumulative reward. In particular, the agent acts on instances of the environment, called states, by taking actions following a certain policy π , that determine the transition to another state. We denote by \mathcal{S} the set of states of the environment, $\mathcal{A}(s)$ the set of actions that the agent can take in state s and by $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ the reward function. A reinforcement learning problem may be described

by a (finite) Markov decision process (MDP) $(s_t, a_t, r_{t+1})_{t \in [0, T-1]}$, where $s_t \in \mathcal{S}$, $a_t \in \mathcal{A}(s_t)$ are the state and action at time t , while $r_{t+1} := r(s_t, a_t)$ is the reward received after action a_t is performed. The time t runs from 0 to T , the time step at which the episode ends. Since the process is Markov, a transition to the next state depends only on the previous state. The goal of the agent is to find an optimal way of behaving, i.e., a policy π that maximizes the discounted cumulative reward, or return,

$$R_t = \sum_{k=0}^{T-t-1} \gamma^k r_{t+1+k}, \quad (1)$$

where $\gamma \in [0, 1]$ is a constant called the discount factor. The value of the discount factor indicates how far in time we take the rewards into account: if it is close to 0, then we have a “myopic” vision of the process, while if it is close to 1 then also actions that happen farther in time are relevant in the computation of the return.

In practice it is very difficult to solve an MDP exactly. For example, very often the transition probabilities are not known and determining them is computationally too expensive even for small scale problems. Thus one often tries to find an approximate solution of the MDP that models the reinforcement learning problem. In particular, we are going to estimate the probabilities by using a deep neural network. In this case, the policy will be denoted by π_θ , where θ denotes the parameters of deep the neural network.

We make use of a popular policy gradient approach called synchronous advantage actor-critic (A2C) [49] to find an approximate solution to the graph partitioning problem. A2C combines the standard REINFORCE algorithm [49], in which updates are made in the direction $\nabla_\theta \log(\pi_\theta(a_t, s_t))R_t$, with θ the parameters of the approximator, with a baseline $(R_t - b_t(s_t))$, that helps to reduce the variance. So the resulting update takes the form $\nabla_\theta \log(\pi_\theta(a_t, s_t))(R_t - b_t(s_t))$. A baseline that is commonly used is the value function of a state $v(s_t)$ [45, 13], a function of a state that, roughly speaking, measures how good it is to be in that state. the term $(R_t - v(s_t))$ is called the advantage. In this model, policy (actor) and value (critic) learning are highly intertwined, resulting in a lower variance and better stability of the approximation. The loss function to be minimized is then

$$L = - \sum_{t=0}^{T-1} \log \pi_\theta(a_t | s_t) (R_t - v(s_t)). \quad (2)$$

In our case, the agent will be modeled by a two-headed deep neural network, with parameter tensor θ , that takes as input a state s and returns a tensor of probabilities $\pi_\theta(\cdot | s)$ for the actions, and a scalar value $v_\theta(s)$ for the value function. Here we use $v_\theta(s)$ to approximate the true value function $v(s_t)$. Hence, the loss function becomes

$$L(\theta) = - \sum_{t=0}^{T-1} \log \pi_\theta(a_t | s_t) (R_t - v_\theta(s_t)) + \alpha \sum_{t=0}^{T-1} (R_t - v_\theta(s_t))^2, \quad (3)$$

where we added the weighted “critic loss” $\alpha \sum_{t=0}^{T-1} (R_t - v_\theta(s_t))^2$ to stabilize the training, with $\alpha \in (0, 1]$.

Figure 3 illustrates the A2C procedure. Given a state, the actor determines what action to take, which will generate a new state. Meanwhile, the critic branch of the network computes the value $v_\theta(s_t)$ of the state, which is used to compute the advantage $(R_t - v_\theta(s_t))$ of the future state s_t . This advantage is then used to reinforce the chosen action a_t . Algorithm 1 illustrates the A2C weight update procedure.

We set the coefficient α to 0.1, in order to let the actor learn faster than the critic. The discount factor γ is set to 0.9, since we are interested in the long term return, as opposed to finding a strictly greedy approach. For the training we use the Adam optimizer [30] with learning rate δ set to 10^{-3} .

Algorithm 2 shows the neural network used for all experiments, which relies on the SAGEConv graph convolutional layer, described in [21]. The SAGE layer implements the relation

$$F'_i \leftarrow F_i W_1 + (\text{mean}_{j \in \mathcal{N}(i)} F_j) W_2, \quad (4)$$

where F is the feature tensor, W_1 and W_2 are weight matrices and $\mathcal{N}(i)$ is the neighborhood of node i . The number of rows in W_1 and W_2 is determined by the number of features in F . In Algorithm 2, the output dimensions are not explicitly mentioned. In the experiments, the number of output channels for the layers is set equal to the number of input channels. The network, as used with 5 input features in Section 3 (finding an edge separator) has 182 tunable parameters. For computing a vertex separator, Section 4, 7 input features




Figure 3: Illustration of the advantage actor-critic (A2C) reinforcement learning approach. Given a state S_t , the actor determines what action to take, generating a new state. Meanwhile, the critic branch of the network estimates the value ($v(S_t)$) of the state, which is used to compute the advantage ($R_t - v(S_t)$) of the future state S_t . This advantage is then used to reinforce the chosen action a_t . This figure is based on an illustration from [50].

Algorithm 1 Update model parameters θ using advantage actor critic (A2C).

Input: model parameters θ ,

rewards r_i , log-probabilities $\log(\pi_i)$ and values v_i for $i \in [1, t]$,

learning rate δ , coefficient α

- 1: **procedure** UPDATE_MODEL_PARAMETERS_A2C($\theta, r, \log(\pi), v$)
 - 2: $R \leftarrow [r_1, r_2 + \gamma r_1, \dots, r_t + \gamma r_{t-1} + \dots + \gamma^{t-1} r_1]$ ▷ compute returns, Eq. 1
 - 3: $R \leftarrow \text{normalize}(R)$
 - 4: ▷ combine actor and critic loss, sum over steps taken, Eq. 3
 - 5: $L \leftarrow -\sum_1^t \log(\pi)(R - v) + \alpha \sum_1^t (v - R)^2$
 - 6: $\theta \leftarrow \theta - \delta \cdot \partial L / \partial \theta$ ▷ back-propagation and step with the optimizer
 - 7: **end procedure**
-

Algorithm 2 Forward evaluation of the agent’s neural network.

Input: graph $G(V, E)$, feature tensor F , I_{mask} list of features with masks for nodes to exclude

Output: actor (action (log-)probabilities), critic (scalar value, only returned in training mode)

- 1: **function** AGENT(G, F, I_{mask})
 - 2: $\text{mask} \leftarrow \text{where}(\bigcup_{i \in I_{\text{mask}}} (F(:, i) \neq 0))$
 - 3: $F \leftarrow \text{Tanh}(\text{SAGEConv}(F))$ ▷ convolutional layers common to actor and critic
 - 4: $F \leftarrow \text{Tanh}(\text{SAGEConv}(F))$
 - 5: $\text{actor} \leftarrow \text{SAGEConv}(F)$ ▷ actor branch
 - 6: $\text{actor}(\text{mask}) \leftarrow -\infty$ ▷ exclude certain actions
 - 7: $\text{actor} \leftarrow \text{log_softmax}(\text{actor})$
 - 8: **if not training then**
 - 9: **return** actor ▷ critic is not required in eval mode
 - 10: **end if**
 - 11: $\text{critic} \leftarrow \text{detach}(F)$ ▷ critic branch
 - 12: $\text{critic} \leftarrow \text{Tanh}(\text{SAGEConv}(\text{critic}))$
 - 13: $\text{critic} \leftarrow \text{Linear}(\text{critic})$ ▷ linear layer
 - 14: $\text{critic} \leftarrow \text{Tanh}(\text{global_mean_pool}(\text{critic}))$ ▷ reduce critic to a single scalar $\in (-1, 1)$
 - 15: **return** actor, critic
 - 16: **end function**
-

are used and the network has 338 parameters. The hyperbolic tangent is used for the nonlinear activation function. The network has two branches, the actor and the critic. The actor contains (log-)probabilities for the possible actions, while the critic predicts the value of the current state. Note that the critic is only required during training. The input to the critic branch of the network, F , is detached in Line 11 so that the common layers are only updated by the actor loss. In Line 14, the critic is reduced to a single scalar value, which, due to the hyperbolic tangent activation function belongs to the interval $(-1, 1)$.

The possible actions are each of the nodes in the graph. What each of these actions means will be discussed in the subsequent sections. However, depending on the context and the current state, certain nodes should not be chosen as they would not lead to valid actions. Therefore, some inputs to the log_softmax layer are set to $-\infty$, so these nodes will never be selected as actions, see Line 2 and Line 6.

We implement Algorithm 2 using PyTorch geometric [17]. To speed up the training process, we run the A2C training using multiple workers with a single shared model, which is called distributed A2C. The workers are created using the PyTorch multiprocessing API.

3 Finding a Minimal Edge Separator

Given a graph $G = (V, E)$, partitioned as $V = V_A \cup V_B$ with $V_A \cap V_B = \emptyset$, its cut is defined as

$$\text{cut}(G) = \sum_{\substack{(v,w) \in E \\ v \in V_A, w \in V_B}} 1, \tag{5}$$

which is simply the number of edges between partitions V_A and V_B . (In the remainder of the paper, when we say $G = (V, E)$ is partitioned into V_A and V_B , it is assumed that $V = V_A \cup V_B$ and $V_A \cap V_B = \emptyset$ even though we will not state it explicitly.) Our goal is to minimize the cut while keeping the two partitions balanced. A popular objective is to minimize the normalized cut, which is defined as

$$\text{NC}(G) = \text{cut}(G) \left(\frac{1}{\text{vol}(V_A)} + \frac{1}{\text{vol}(V_B)} \right), \tag{6}$$

with the volume of a partition defined as

$$\text{vol}(V_A) = \sum_{v \in V_A} \text{deg}(v). \tag{7}$$

For a pictorial example of the cut and the normalized cut see Figure 1.

In Algorithm 3, we present an algorithm to find a partition $V = V_A \cup V_B$ for the graph $G(V, E)$ that approximately minimizes the normalized cut. The outline of the algorithm is as follows. If the graph G is large enough, the algorithm constructs a coarse representation G^C of the graph, and Algorithm 3 is applied recursively on this coarser graph G^C . To stop the recursion, when the graph is small enough, i.e., when $|V| \leq n_{\min}$, the graph is partitioned directly using for instance the METIS graph partitioner (Line 3), or a separately trained reinforcement learning-based graph partitioning algorithm, see Section 3.2. The result of the recursive call is a partitioning $V^C = V_A^C \cup V_B^C$ of the coarser graph, which is then interpolated back (Line 7) to the finer graph G . The resulting partitioning, which should already be relatively good, is then refined using a deep reinforcement learning approach (Line 15 - Line 27). This multilevel approach is sketched in Figure 2. However, to speed up the refinement process, a subgraph G^{sub} is constructed (Line 10) with the nodes in G that are within a small number of hops from the cut. The cut is then only refined within this subgraph.

For the coarsening step, we use the graclus [5] graph clustering code, which is based on heavy edge matching and groups nodes in clusters of size 2, with a small number of unmatched nodes leading to clusters of size 1. These clusters define the nodes for the coarse graph G^C . The coarsening rate is typically slightly less than 2, leading to $\sim \log |V|$ recursion levels. For the interpolation of the coarse partitioning back to the finer graph, nodes in G which correspond to coarse nodes in V_A^C , or V_B^C , are all assigned to V_A , or V_B respectively, in the fine graph (Line 7). In Algorithm 3, I^C denotes the mapping from coarse to fine nodes. Note that the first step in the graclus clustering algorithm is a random permutation of the nodes, which

Algorithm 3 Computing an edge separator using deep reinforcement learning. This illustrates both the training and evaluation on a single graph.

Input: graph $G(V, E)$

Output: partitions V_A and V_B , such that $V = V_A \cup V_B$

```

1: function EDGE_SEPARATOR( $G$ )
2:   if  $|V| < n_{\min}$  then                                     ▷ end the recursion
3:     return metis_partition( $G$ ) or edge_separator_coarse( $G$ )    ▷ See Section 3.2
4:   end if
5:    $G^C, I^C \leftarrow$  coarsen( $G$ )                               ▷ get coarse graph and interpolation info
6:    $V_A^C, V_B^C \leftarrow$  metis_partition( $G^C$ ) if training else edge_separator( $G^C$ )    ▷ recursion
7:    $V_A, V_B \leftarrow V_A^C(I^C), V_B^C(I^C)$                    ▷ interpolate  $V_A$  and  $V_B$  from coarse to fine
8:    $V_A^0, V_B^0 \leftarrow V_A, V_B$                              ▷ keep a copy
9:    $V_C \leftarrow \{v\} \cup \{w\}, \forall v, w : \exists e_{v,w}$  with  $v \in V_A, w \in V_B$     ▷ find nodes around the cut
10:   $G^{\text{sub}} \leftarrow$  k_hop_subgraph( $G, V_C, k_{\text{hops}}$ )          ▷ subgraph with all nodes at most  $k_{\text{hops}}$  from  $V_C$ 
11:                                          ▷ construct feature tensor
12:   $F(v) \leftarrow [v \in V_A, v \in V_B, v \in \partial G^{\text{sub}}, \text{vol}(V_A)/\text{vol}(V), \text{vol}(V_B)/\text{vol}(V)]$ ,  $\forall v \in G^{\text{sub}}$ 
13:   $c \leftarrow$  cut( $G, V_A, V_B$ )                                 ▷ compute the cut size, Eq. 5
14:   $C \leftarrow$  NC( $G, V_A, V_B$ )                               ▷ compute normalized cut, Eq. 6
15:  for  $t \leftarrow 1$  to  $c$  do
16:    if training then
17:      policy $_t$ , value $_t \leftarrow$  agent( $G^{\text{sub}}, F, 2$ )        ▷ forward evaluation of the agent, see Algorithm 2
18:       $a_t \leftarrow$  categorical_sample_logits(policy $_t$ )      ▷ pick action
19:    else
20:      policy $_t \leftarrow$  agent( $G^{\text{sub}}, F, 2$ )                ▷ forward evaluation of the agent, see Algorithm 2
21:       $a_t \leftarrow$  argmax(policy $_t$ )                        ▷ pick action
22:    end if
23:     $V_A, V_B \leftarrow V_A \setminus a_t, V_B \cup a_t$  if  $a_t \in V_A$  else  $V_A \cup a_t, V_B \setminus a_t$     ▷ move  $a_t$  from  $V_A/V_B$  to  $V_B/V_A$ 
24:     $F(v) \leftarrow [v \in V_A, v \in V_B, v \in \partial G^{\text{sub}}, \text{vol}(V_A)/\text{vol}(V), \text{vol}(V_B)/\text{vol}(V)]$ ,  $\forall v \in G^{\text{sub}}$     ▷ update features
25:     $C_{\text{old}}, C \leftarrow C, \text{NC}(G, V_A, V_B)$               ▷ compute normalized cut, Eq. 6
26:     $r_t \leftarrow C_{\text{old}} - C$                                ▷ compute reward
27:  end for
28:  if training then
29:    update_model_parameters_A2C(agent,  $r$ , policy, value)
30:  else
31:     $V_A, V_B \leftarrow V_A^0, V_B^0$ 
32:    for  $t \leftarrow 1$  to argmax( $r$ ) do
33:       $V_A, V_B \leftarrow V_A \setminus a_t, V_B \cup a_t$  if  $a_t \in V_A$  else  $V_A \cup a_t, V_B \setminus a_t$ 
34:    end for
35:  end if
36:  return  $V_A, V_B$ 
37: end function

```

means that the graph coarsening phase is not deterministic. Since the partition quality depends also on the graph coarsening, for each graph the algorithm is repeated 3 times and the best partitioned graph is kept.

The procedure in Algorithm 3 handles both the training and the evaluation for a single graph. During the training Algorithm 3 is called for each graph in the training dataset, and this is repeated for multiple epochs.

3.1 Refinement of the Cut using Deep Reinforcement Learning

Line 15 to Line 27 of Algorithm 3 show a single episode of the deep reinforcement learning algorithm to refine the interpolated partition. Let c denote the size of the cut of the interpolated partitioning, computed in Line 13. Since it is assumed that the algorithm achieved a high quality partitioning on the coarser problem, we expect to only require $\mathcal{O}(c)$ steps to refine the cut on the finer graph in order to overcome imperfections introduced by the interpolation procedure. Therefore we set the episode length to c (Line 15). From the experiments we observe that taking more than c steps typically does not further improve the partitioning. In every step of the episode one node from G^{sub} is selected and used to perform an action. During training, this node is selected by sampling (Line 18) from the agent’s policy. The policy contains log-probabilities for each of the nodes in G^{sub} . The policy corresponds to the actor output from the agent’s neural network (see Algorithm 2) applied to the graph G^{sub} and the corresponding node feature tensor F . The feature tensor is discussed in more detail below. During evaluation, the node with highest probability is selected (Line 21). When a node a_t is selected, an action is taken (Line 23). The action at step t is also denoted as a_t , and we use a to refer to the vector with all actions from step 1 up to the current step t . If the selected node a_t is in V_A , then it is moved to V_B , alternatively, if a_t was in V_B , it is moved to V_A . At this point the normalized cut is computed for the new state, and the reward r_{t+1} at step t is defined as the difference between the previous and the new normalized cut (Line 26). Note that if a_t is moved from V_A to V_B , the volumes $\text{vol}(V_A)$ and $\text{vol}(V_B)$ can simply be updated by subtracting and adding $\text{deg}(a_t)$ respectively. Likewise, the cut c can be updated cheaply by only considering the neighbors of a_t .

The Feature Tensor At Line 12 the feature tensor F is constructed, with 5 features per node in the graph G^{sub} . The first two features are a one-hot encoding of the partition the node belongs to, with $[1, 0]$ referring to V_A and $[0, 1]$ to V_B . Let ∂G^{sub} denote the boundary of the subgraph G^{sub} around the cut, i.e., the nodes in G^{sub} which are connected to nodes in $G \setminus G^{\text{sub}}$. These nodes have edges which are not part of G^{sub} and are hence not seen by the agent. Therefore, these nodes should not be selected as actions. Thus, the next feature denotes whether (1) or not (0) a node is in ∂G^{sub} . Nodes with this feature set to 1 will never be selected, since in the agent’s neural network their input to the softmax layer is set to minus infinity, see Line 2, 6 in Algorithm 2. In our case, the softmax layer is defined as a layer σ that takes as input a $(m, 1)$ -tensor z and returns

$$\sigma(z)_i = \frac{\exp(z_i)}{\sum_{j=1}^m \exp(z_j)}, \quad (8)$$

where m is the number of nodes of the input graph G^{sub} . The final two features are the normalized volumes of the two partitions: $\text{vol}(V_A)/\text{vol}(V)$ and $\text{vol}(V_B)/\text{vol}(V)$. These volumes play an important role in the reward function, and they cannot be determined from G^{sub} and the other features alone. The feature tensor is updated in every step of the episode, after an action is taken, on Line 24. Table 1 summarizes the different features used for each deep reinforcement learning model.

In Algorithm 3, r is used to denote the vector of all rewards from the first steps up to the current step t . Similarly, policy_t refers to the log-probabilities at step t and policy contains all these log-probabilities stacked together. In Algorithm 3, to simplify the notation, the network parameters are updated only once, after the entire episode is finished. The update of the model parameters is done in Line 29, with a call to the A2C update procedure Algorithm 1. However, in practice, the network parameters θ are updated after a fixed number of steps in the episode, as well as at the end of the episode.

In evaluation mode, only the actions that actually contribute to the peak cumulative reward are applied to the graph. Therefore, a copy is made of the initial partitioning V_A, V_B before the start of the episode,

Algorithm	Feature tensor			
	Partition info		Binary mask	Imbalance info
Edge-cut refining	$v \in V_A$	$v \in V_B$	$v \in \partial G^{\text{sub}}$	$\frac{\text{vol}(V_A)}{\text{vol}(V)}$ $\frac{\text{vol}(V_B)}{\text{vol}(V)}$
Edge-cut coarsest graph	$v \in V_A$	$v \in V_B$		
Vertex separator	$v \in V_A$	$v \in V_B$	$v \in V_S$ $v \in \partial G^{\text{sub}}$	$v \in V_S^{\text{min}}$ $ V_A / V $ $ V_B / V $

Table 1: Feature tensors for each considered model. On the left we list the studied models: edge-cut refining, coarsest graph partitioning and vertex separator. On the right, each feature is positioned in a specific column according if it is related to the partition information, it serves as a mask or it is an imbalance information. Binary mask features exclude the node from being picked during the DRL process.




Figure 4: Representation of the algorithm to partition the coarsest graph. At the first step all nodes are in one partition (blue) except for one node with minimum degree (the red one). Then, at each step the agent picks a blue node and turns it into a red one, until there are exactly the same number of blue and red vertices.

see Line 8. Line 32 to Line 34 apply the actions from the episode that lead to the peak cumulative reward obtained during the episode.

3.2 Partitioning the Coarsest Graph

In order to partition the coarsest graphs (see Line 3 in Algorithm 3) we train a separate actor-critic reinforcement learning agent. Given a graph G with n vertices, we assign all of its nodes to partition B , except for one, one of the nodes with smallest degree, which is assigned to partition A . By choosing the node with smallest degree the initial normalized cut is as small as possible. In this case, each node v in B has feature $F(v) = [1, 0]$, while each node v in A has feature $F(v) = [0, 1]$, see Table 1. During the training, the agent picks a node in partition B , according to the output probabilities of the deep neural network, and moves it to partition A . This corresponds to changing the feature $F(v)$ of the chosen node from $[1, 0]$ to $[0, 1]$. The process is repeated until the two partitions reach the same cardinality: $|A| = |B|$ if n is even or $|A| = |B| + 1$ if n is odd. During evaluation we allow an imbalance of $\iota = 1\%$ of n between the two partitions. Algorithm 4 shows the training and the evaluation process explained above. Figure 4 illustrates the above algorithm on a toy example.

The deep neural network used for this task is slightly different than the one (Algorithm 2) for the refinement phase. An important ingredient is the attention mechanism on graphs, which is implemented by graph attention (GAT) [47] layers. The neural network has the actor and the critic branches after 4 common GAT layers. For the actor, a few dense layers are used, while for the critic branch we use a global pooling layer [37], which allows to get a scalar value. In this case, the nodes that have already been picked, i.e., the nodes with feature $[0, 1]$, are masked. All convolutional layers have 10 units, while all linear layers have 5 units. The global pooling layer includes a dense neural network with 2 linear layers, with 5 and 1 units respectively, and Tanh as activation function in between. Algorithm 5 shows the structure of the neural network.

3.3 Algorithm Complexity

For a graph with $n = k^2$ nodes resulting from the spatial discretization of a regular, square, $k \times k$ two-dimensional problem, the size of the cut will be $c = \mathcal{O}(n^{1/2})$. Likewise, for a three-dimensional problem,

Algorithm 4 Computing an edge separator using deep reinforcement learning on the coarsest graph. This illustrates both the training and evaluation on a single graph.

Input: graph $G(V, E)$, with $|V| = n$. Imbalance factor ι .

Output: partitions V_A and V_B , such that $V = V_A \cup V_B$

```

1: function EDGE_SEPARATOR_COARSE( $G$ )
2:    $V_A^0, V_B^0 \leftarrow V_A, V_B$  ▷ keep a copy
3:    $C \leftarrow \text{NC}(G, V_A, V_B)$  ▷ compute normalized cut, Eq. 6
4:   for  $t \leftarrow 1$  to  $n/2 - 1$  do
5:     if training then
6:       policy $_t$ , value $_t \leftarrow$  agent( $G, F, 2$ ) ▷ forward evaluation of the agent, see Algorithm 5
7:        $a_t \leftarrow$  categorical_sample_logits(policy $_t$ ) ▷ pick action
8:     else
9:       policy $_t \leftarrow$  agent( $G, F, 2$ ) ▷ forward evaluation of the agent, see Algorithm 5
10:       $a_t \leftarrow$  argmax(policy $_t$ ) ▷ pick action
11:    end if
12:     $V_A, V_B \leftarrow V_B \setminus a_t, V_A \cup a_t$  ▷ move  $a_t$  from  $V_A/V_B$ 
13:     $F(a_t) \leftarrow [0, 1]$  ▷ update features
14:     $C_{\text{old}}, C \leftarrow C, \text{NC}(G, V_A, V_B)$  ▷ compute normalized cut, Eq. 6
15:     $r_t \leftarrow C_{\text{old}} - C$  ▷ compute reward
16:  end for
17:  if training then
18:    update_model_parameters_A2C(agent,  $r$ , policy, value)
19:  else
20:     $V_A, V_B \leftarrow V_A^0, V_B^0$ 
21:    for  $t \leftarrow 1$  to  $n/2 - 1 + \iota n/100$  do
22:       $V_A, V_B \leftarrow V_B \setminus a_t, V_A \cup a_t$ 
23:    end for
24:  end if
25:  return  $V_A, V_B$ 
26: end function

```

Algorithm 5 Forward evaluation of the agent’s neural network on the coarsest graph.

Input: graph $G(V, E)$, feature tensor F , I_{mask} list of features with masks for nodes to exclude

Output: actor (action (log-)probabilities), critic (scalar value, only returned in training mode)

```

1: function AGENT( $G, F, I_{\text{mask}}$ )
2:   mask  $\leftarrow$  where(  $\bigcup_{i \in I_{\text{mask}}} (F(:, i) \neq 0)$ )
3:    $F \leftarrow$  Tanh(GATConv( $F$ )) ▷ convolutional layers common to actor and critic
4:    $F \leftarrow$  Tanh(GATConv( $F$ ))
5:    $F \leftarrow$  Tanh(GATConv( $F$ ))
6:    $F \leftarrow$  Tanh(GATConv( $F$ ))
7:    $F \leftarrow$  Tanh(Linear( $F$ )) ▷ dense layers common to actor and critic
8:    $F \leftarrow$  Tanh(Linear( $F$ ))
9:   actor  $\leftarrow$  Tanh(Linear( $F$ )) ▷ actor branch
10:  actor  $\leftarrow$  Tanh(Linear(actor))
11:  actor(mask)  $\leftarrow -\infty$  ▷ exclude certain actions
12:  actor  $\leftarrow$  log_softmax(actor)
13:  if not training then
14:    return actor ▷ critic is not required in eval mode
15:  end if
16:  critic  $\leftarrow$  GlobalAttention( $F$ ) ▷ critic branch
17:  critic  $\leftarrow$  Tanh(Linear(critic))
18:  critic  $\leftarrow$  Linear(critic)
19:  return actor, critic
20: end function

```

the cut will be $c = \mathcal{O}(n^{2/3})$, i.e., a plane through the domain. In Algorithm 3, the evaluation of the neural network (Line 20) on G^{sub} takes $\mathcal{O}(c)$ computations, since G^{sub} has $\mathcal{O}(c)$ nodes and the network has a fixed, small number of convolutional layers. The refinement procedure takes at most $\mathcal{O}(c)$ steps, so the total cost for the refinement at the finest level is $\mathcal{O}(c^2)$ or $\mathcal{O}(n)$ in 2D and $\mathcal{O}(n^{4/3})$ in 3D. Since the coarsening rate is close to 2, the total cost for Algorithm 3 is

$$\text{cost}_{2D} = \sum_{i=1}^{\log n} \frac{n}{2^{i-1}} = \mathcal{O}(n), \quad \text{cost}_{3D} = \sum_{i=1}^{\log n} \left(\frac{n}{2^{i-1}} \right)^{4/3} = \mathcal{O}(n^{4/3}), \quad (9)$$

for 2D and 3D respectively.

3.4 Experimental Evaluation

This section discusses experiments with the proposed DRL partitioning algorithm. We consider two types of graphs: triangulations and graphs from a variety of different applications, all from spatial discretizations. For the first class, training is performed on a set of Delaunay triangulations and the resulting algorithm is tested on other triangulations, constructed using finite element modeling. In the second case, we train on graphs from 2D and 3D discretizations from a matrix collection commonly used as benchmark.

Delaunay Triangulations The graphs in this test correspond to planar Delaunay triangulations from points randomly generated in the unit square, see for example Figure 10a. The training dataset contains roughly N^{train} graphs with $n^{\text{train}} \in [n_{\text{min}}^{\text{train}}, n_{\text{max}}^{\text{train}}]$ nodes. This set is constructed as follows. A random Delaunay graph with n^{train} (sampled uniformly from $[n_{\text{min}}^{\text{train}}, n_{\text{max}}^{\text{train}}]$) nodes is generated and added to the dataset. Then this last graph is coarsened, and the coarsened graph is added to the dataset. As long as the coarser graph has more than $n_{\text{min}}^{\text{train}}$ nodes, this is repeated by coarsening the coarser graph again and adding the coarser version to the dataset as well. These operations are repeated until the dataset has N^{train} elements. The graphs in the test set are simply randomly generated Delaunay triangulations, without the coarsenings. The parameter $n_{\text{min}}^{\text{train}}$ is also used to stop the recursive coarsening during the evaluation phase, see Line 3 in Algorithm 3.

	Train	Test ①	Test ②	Test ③	Test ④	
N	Delaunay	10000	20	127	181	172
	GradedL	-	16	26	6	-
	Hole3	-	11	20	13	5
	Hole6	-	11	20	13	5
n_{\min}	100	100	5000	30000	60000	
n_{\max}	5000	5000	30000	60000	90000	

Table 2: Parameters describing the datasets. The graphs in the Delaunay dataset are constructed from $n_{\min} < n \leq n_{\max}$ random points in $[0, 1]^2$. The GradedL, Hole3 and Hole6 dataset are finite element triangulations of 3 different geometries, for multiple levels of mesh refinement. The test sets are split in 4 separate ranges.

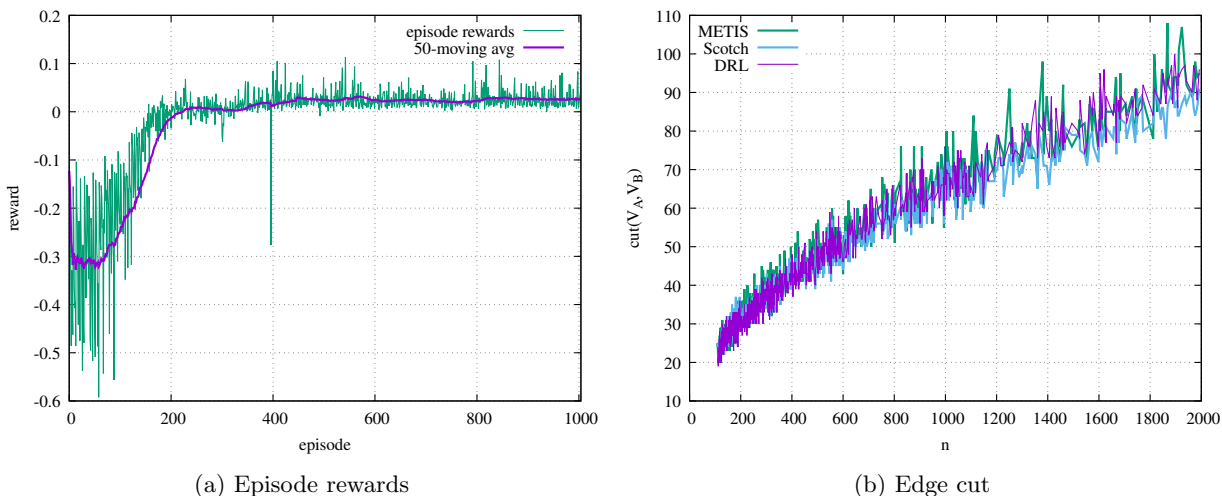


Figure 5: Training on dataset with Delaunay graphs from random points in $[0, 1]^2$. Figure 5a: Total rewards of 1000 episodes. The final rewards are consistently positive. Figure 5b: Cut sizes of all graphs in the training dataset, sorted based on number of nodes, as computed with the proposed DRL method, METIS and SCOTCH. For these 2D graphs, the cut size scales as \sqrt{n} .

We consider four different tests, with different parameters as detailed in Table 2.

Figure 5a shows the cumulative rewards for episodes on each of the graphs in the training dataset. Figure 5b compares the cut sizes on the training graphs for the proposed method, METIS and SCOTCH. This also illustrates that for planar problems on a square domain, as expected, the cut size scales as \sqrt{n} .

Figure 6 compares the normalized cut, the cut and the balance on the Delaunay testing dataset obtained by our model, METIS and SCOTCH. The balance is measured by $\max\left\{\frac{\text{vol}(A)}{\text{vol}(B)}, \frac{\text{vol}(B)}{\text{vol}(A)}\right\}$. We see that, in every range of nodes, the normalized cut is very close to the one produced by METIS and SCOTCH and the partitions are balanced as well. The edge cut results are slightly higher for graphs having from 60,000 to 90,000 nodes.

Figure 10 illustrates the partitioning of one example of a Delaunay graph. Figure 10b shows the sub-graph G^{sub} , consisting of all nodes at distance at most 3 hops from the edge cut.

Finite Element Triangulations Figure 11 illustrates the finite element triangulations we use for testing. We consider three different meshes – GradedL, Hole3 and Hole6 – each with multiple levels of refinement. Figures 7, 8 and 9 show the normalized cut, the cut size and the balance for these datasets, and the comparison with the ones obtained with METIS and SCOTCH. Notice that partitioning quality is close to that of METIS or SCOTCH. Recall that these graphs have a very different sparsity pattern with respect to the ones in the training dataset (Delaunay graphs), so this shows that the agent is able to generalize well on

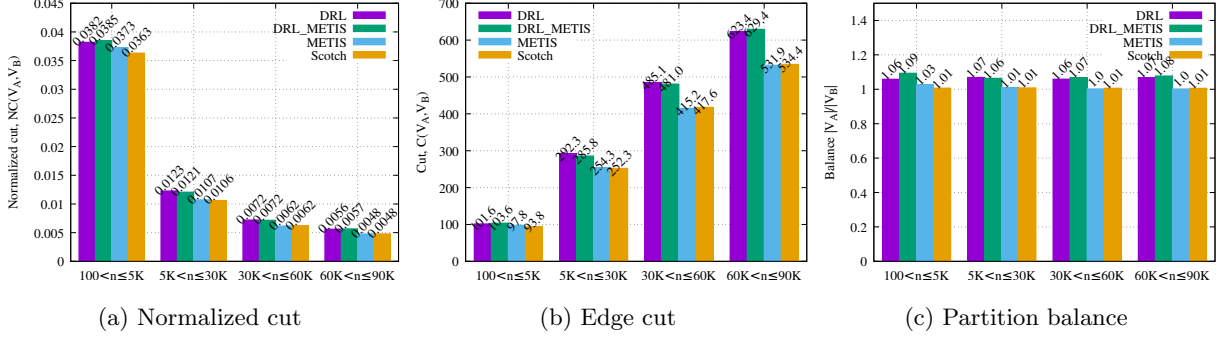


Figure 6: Evaluation of the partitioning algorithms on the Delaunay testing sets. DRL and DRL_METIS refer to Algorithm 3. DRL_METIS uses METIS on the coarser level, while DRL uses reinforcement learning partitioning on the coarsest level as well.




Figure 7: Evaluation of the partitioners on the GradedL triangulation dataset. Training was performed on Delaunay graphs with $100 < n \leq 5000$ nodes.

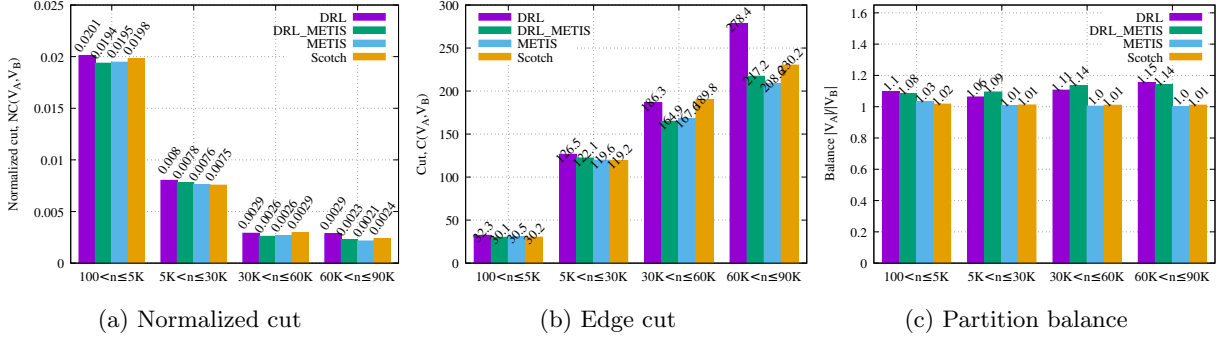


Figure 8: Evaluation of the partitioners on the Hole3 (see Figure 11b) triangulation dataset. Training was performed on Delaunay graphs with $100 < n \leq 5000$ nodes.




Figure 9: Evaluation of the partitioners on the Hole6 (see Figure 11c) triangulation dataset. Training was performed on Delaunay graphs with $100 < n \leq 5000$ nodes.

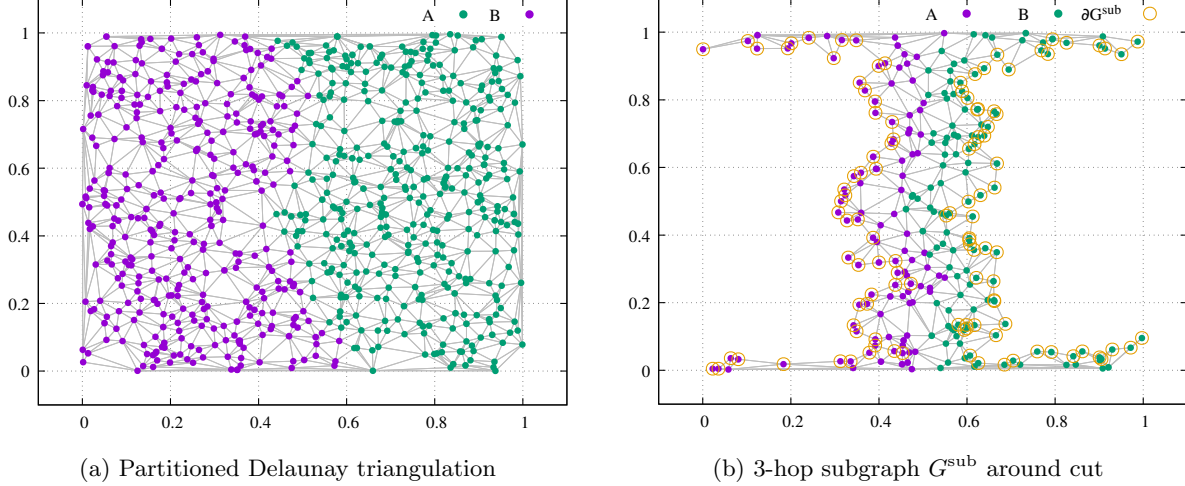


Figure 10: Figure 10a shows a Delaunay triangulation from 750 random points in $[0, 1]^2$, partitioned using METIS. Call V_A the partition including the purple nodes and V_B the partition including the green nodes. Figure 10b illustrates the 3-hop subgraph around the edge cut. The purple circular nodes have features $[1, 0, 0, \text{vol}(V_A)/\text{vol}(V), \text{vol}(V_B)/\text{vol}(V)]$ while the green ones have features $[0, 1, 0, \text{vol}(V_A)/\text{vol}(V)]$. The nodes bounded by an orange circle belong to the boundary. More precisely, the purple orange-bounded nodes have features $[1, 0, 1, \text{vol}(V_A)/\text{vol}(V), \text{vol}(V_B)/\text{vol}(V)]$, while the green orange-bounded nodes have features $[0, 1, 1, \text{vol}(V_A)/\text{vol}(V), \text{vol}(V_B)/\text{vol}(V)]$.




Figure 11: Illustrations for the different finite element triangulations, Figure 11a GradedL after 15 refinements, Figure 11b Hole3 after 10 refinements, and Figure 11c Hole6 after 10 refinements.




Figure 12: Evaluation of the partitioners on the SuiteSparse dataset. Training was performed on graphs from the SuiteSparse collection with $100 < n \leq 5000$ nodes.

unseen planar triangulations.

Sparse Matrices from the SuiteSparse Collection We also tested our model on a number of 2D and 3D discretizations from the SuiteSparse Matrix Collection [33]. We trained a separate agent, with a deep neural network having the same structure as the one used for the Delaunay triangulations. The training dataset is built in the same fashion as we did for the training on Delaunay triangulations. More precisely, we pick a (fully connected) 2D/3D discretization having number of vertices in $(100, 5000]$ and we add it to the dataset. Then we coarsen it until the coarsest graph has less than 100 nodes and we add all the coarser graphs to the dataset. These operations are repeated until the dataset has 10,000 elements. It is important to stress that training a separate agent for these kind of graphs is necessary, since they have a great variety of sparsity patterns. The agent trained on Delaunay triangulations may fail to generalize correctly, in particular in providing balanced partitions.

In order to see if the model is able to generalize to unseen graphs, we tested it on larger 2D/3D discretizations, with between 5,000 and 90,000 nodes (again, for testing we do not include coarser graphs). Figure 12 shows the normalized cut, the cut size and the balance for these graphs. The agent is able to generalize to these graphs, producing a lower cut and a lower normalized cut. The partitions are slightly unbalanced for graphs in the mid range of nodes.




Figure 13: Example of a vertex separator. The vertices in the separator are depicted in green, while the other two partitions are depicted in red and blue. Note that the subgraph obtained by removing the vertices in the separator has exactly two connected components. In this case, the two components have approximately the same cardinality.

4 Finding a Minimal Vertex Separator

A vertex separator is a set of nodes that, if removed from the graph, would split the graph in two unconnected sub-graphs, see Figure 13. Hence, the graph is partitioned in three sub-graphs $G_A = (V_A, E_A)$, $G_B = (V_B, E_B)$ and $G_S = (V_S, E_S)$, such that $V = V_A \cup V_B \cup V_S$ and there are no edges connecting V_A and V_B . The aim is to minimize $|V_S|$ while keeping the two other graphs V_A and V_B balanced. While the normalized cut (Eq. 6), which was used for graph bisection as discussed in Section 3, is based on the volumes of the partitions (Eq. 7), we will now try to balance the cardinalities $|V_A|$ and $|V_B|$. More precisely, we attempt to minimize

$$\text{NS}(G) = |V_S| \left(\frac{1}{|V_A|} + \frac{1}{|V_B|} \right), \quad (10)$$

a measure for the normalized separator. Note that a vertex separator can be computed using a number of heuristics [35, 23], or can be constructed from an edge separator using a minimum cover approach [41, 15].

Algorithm 6 computes, for a given graph $G(V, E)$, a vertex separator V_S that approximately minimizes Eq. 10. Algorithm 6 follows the same structure as Algorithm 3, with some significant differences. Most importantly, the reward used in Algorithm 6 is based on Eq. 10. In Line 3, to stop the recursion, now METIS is called to compute a vertex separator instead of an edge separator.

As before, a separator (now a vertex separator) is computed first on the coarser graph and then interpolated back to the finer graph. Then this separator is refined using deep reinforcement learning applied to a subgraph G^{sub} containing all nodes within a small number of hops from the vertex separator. In every step of the deep reinforcement learning episode a single node a_t is selected. Algorithm 7 illustrates the action that is taken for a selected node a_t . If node a_t is part of V_A , it is simply removed from partition V_A and added to the separator V_S (Line 3). Likewise, if a_t was in V_B , then node a_t is removed from V_B and added to V_S (Line 5). The partition volumes can be updated directly. If the selected node a_t is part of the separator, we try to remove it from the separator V_S and move it to one of the partitions V_A or V_B . However, doing so could result in a state where V_S is no longer a valid separator. This would be the case if after removing a_t from the separator – and moving it to either V_A or V_B – there would be a direct edge between V_A and V_B . However, by construction of the feature tensor and the agent’s neural network, such a node will never be selected, as discussed in more detail below. Algorithm 7 thus assumes that it is safe to remove a_t from the separator. To decide whether to move a_t from V_S to V_A or, instead to V_B , it is checked if a_t was already connected to either V_A or V_B . If there is an edge from a_t to any node in V_A , then a_t can be moved to V_A (Line 10). Likewise, if a_t was already connected to V_B it can be moved to V_B (Line 11). If a_t has no edge to either V_A or V_B , then a_t is added to whichever one of V_A or V_B has the smallest cardinality (Line 12). Doing so will always result in a positive reward: the separator becomes smaller and the balance improves.

Since we assume that the vertex separator computed at the coarser graph is of high quality, only a small number of refinement steps will be required. However, since moving the vertex separator by one node requires two actions – first adding a node to the separator, then removing one – the length of an episode is set to $2|V_S^0|$, i.e., twice the size of the interpolated separator.

Algorithm 6 Computing a vertex separator using deep reinforcement learning. Shown here is only the evaluation phase.

Input: graph $G(V, E)$

Output: vertex separator V_S , unconnected components V_A and V_B , such that $V = V_S \cup V_A \cup V_B$

```

1: function VERTEX_SEPARATOR( $G$ )
2:   if  $|V| < n_{\min}$  then
3:     returnmetis_separator( $G$ )
4:   end if
5:    $G^C, I^C \leftarrow$  coarsen( $G$ )                                 $\triangleright$  get coarse graph and interpolation info
6:    $V_S^C, V_A^C, V_B^C \leftarrow$  vertex_separator( $G^C$ )           $\triangleright$  recursive call using coarse graph
7:    $V_S, V_A, V_B \leftarrow V_S^C(I^C), V_A^C(I^C), V_B^C(I^C)$    $\triangleright$  interpolate  $V_S, V_A$  and  $V_B$  from coarse to fine
8:    $V_S^0, V_A^0, V_B^0 \leftarrow V_S, V_A, V_B$                  $\triangleright$  keep a copy
9:    $G^{\text{sub}} \leftarrow$  k_hop_subgraph( $G, V_S, k_{\text{hops}}$ )         $\triangleright$  subgraph with all nodes at most  $k_{\text{hops}}$  from  $V_S$ 
10:   $F(v) \leftarrow [v \in V_A, v \in V_B, v \in V_S, v \in \partial G^{\text{sub}}, v \in V_S^{\min}, |V_A|/|V|, |V_B|/|V|], \forall v \in G^{\text{sub}}$   $\triangleright$  construct feature tensor
11:   $S \leftarrow$  NS( $V_S, V_A, V_B$ )                                $\triangleright$  compute normalized separator, Eq. 10
12:  for  $t \leftarrow 1$  to  $2|V_S^0|$  do
13:    policy $_t \leftarrow$  agent( $G^{\text{sub}}, F, \{2, 3\}$ )              $\triangleright$  forward evaluation of agent's neural network
14:     $a_t \leftarrow$  argmax(policy $_t$ )                            $\triangleright$  pick action
15:    apply_action_vertex_separator( $a_t, G, V_S, V_A, V_B$ )       $\triangleright$  take an action, see Algorithm 7
16:     $\triangleright$  update features
17:     $F(v) \leftarrow [v \in V_A, v \in V_B, v \in V_S, v \in \partial G^{\text{sub}}, v \in V_S^{\min}, |V_A|/|V|, |V_B|/|V|], \forall v \in G^{\text{sub}}$ 
18:     $S_{\text{old}} \leftarrow S$ 
19:     $S \leftarrow$  NS( $V_S, V_A, V_B$ )                              $\triangleright$  compute normalized separator, Eq. 10
20:     $r_t \leftarrow S_{\text{old}} - S$                                 $\triangleright$  compute reward
21:  end for
22:   $V_S, V_A, V_B \leftarrow V_S^0, V_A^0, V_B^0$ 
23:  for  $t \leftarrow 1$  to argmax( $r$ ) do
24:    apply_action_vertex_separator( $a_t, G, V_S, V_A, V_B$ )       $\triangleright$  take an action, see Algorithm 7
25:  end for
26:  return  $V_S, V_A, V_B$ 
27: end function

```

Algorithm 7 Apply an action to a vertex separator. This is used in Algorithm 6 to compute a minimal vertex separator. This assumes node a is not an essential node for the separator, i.e., node a is not connected to both V_A and V_B .

Input: node a , graph $G(V, E)$, partitions V_S, V_A and V_B

Result: node a is moved to a different partition

```

1: procedure APPLY_ACTION_VERTEX_SEPARATOR( $a, G(V, E), V_S, V_A, V_B$ )
2:   if  $a \in V_A$  then
3:      $V_A, V_S \leftarrow V_A \setminus a, V_S \cup a$                                 ▷ move node  $a$  from  $V_A$  to  $V_S$ 
4:   else if  $a \in V_B$  then
5:      $V_B, V_S \leftarrow V_B \setminus a, V_S \cup a$                                 ▷ move node  $a$  from  $V_B$  to  $V_S$ 
6:   else if  $a \in V_S$  then
7:      $V_S \leftarrow V_S \setminus a$ 
8:     if  $\exists e_{a,v} \in E : v \in V_A$  then                                        ▷ node  $a$  is in  $V_S$  and connected to  $V_A$ 
9:        $V_A \leftarrow V_A \cup a$                                                 ▷ move  $a$  to  $V_A$ 
10:    else if  $\exists e_{a,v} \in E : v \in V_B$  then                                    ▷ node  $a$  is in  $V_S$  and connected to  $V_B$ 
11:       $V_B \leftarrow V_B \cup a$                                                 ▷ move  $a$  to  $V_B$ 
12:    else if  $|V_A| \leq |V_B|$  then                                            ▷  $a$  is not connected to  $V_A$  or  $V_B$ 
13:       $V_A \leftarrow V_A \cup a$                                                 ▷ moving  $a$  to  $V_A$  improves the balance
14:    else                                                                    ▷  $|V_B| < |V_A|$ 
15:       $V_B \leftarrow V_B \cup a$                                                 ▷ moving  $a$  to  $V_B$  improves the balance
16:    end if
17:  end if
18: end procedure

```

For brevity, the training is omitted from Algorithm 6, since it can be added similarly to Algorithm 3.

The Feature Tensor The feature tensor used in Algorithm 6 is similar to the one used in Algorithm 3 to compute an edge separator. However, there are two additional features. One extra features comes from the fact that there are now three partitions instead of two, so the one-hot encoding of the partitioning is now $[1, 0, 0]$, $[0, 1, 0]$ and $[0, 0, 1]$ for V_A , V_B and V_S respectively. The next feature again denotes whether a node is part of ∂G^{sub} , in which case it should not be selected as an action. The next features is a binary feature that is set to 1 for the nodes that are essential to the separator, i.e., removing this node from the separator would lead to an invalid state. In Algorithm 6 this set of nodes is denoted as V_S^{min} , and these are all nodes $v \in V_S : \exists e_{v,i}, e_{v,j}$ with $i \in V_A, j \in V_B$, i.e., node v is in the separator and is connected to both V_A and V_B . The final two features are $|V_A|/|V|$ and $|V_B|/|V|$.

4.1 Experimental Evaluation

Delaunay Triangulations We train the algorithm on Delaunay triangulations with between 100 and 1,000 nodes. As for the normalized cut, all intermediate coarsened graphs are included in the dataset, for a total of $\sim 10,000$ training graphs.

Figure 14a shows the normalized separator, Eq. 10, for the Delaunay testing sets corresponding to the parameters as given in Table 2. Figure 14b shows the normalized separator for the Hole6 dataset. For the DRL experiments, METIS is used to find a vertex separator at the coarsest level. Our model produces slightly larger normalized separators than METIS, but the difference is small in all the different tests. Also, the algorithm is able to generalize extremely well also to the unseen triangulations of the Hole6 dataset, and it outperforms METIS for graphs having between 5,000 and 30,000 vertices.

Sparse matrices from the SuiteSparse matrix collection As in the edge cut case, we train our agent on some 2D/3D discretizations from the SuiteSparse matrix collection. The dataset is the same as the one discussed in Section 3.4 for the SuiteSparse experiments, and so is the structure of the deep neural network.

Testing was done on 101 2D/3D discretizations from the SuiteSparse dataset with number of vertices between 5,000 and 90,000. Figure 14c shows the results for this dataset. Also in this case, the normalized

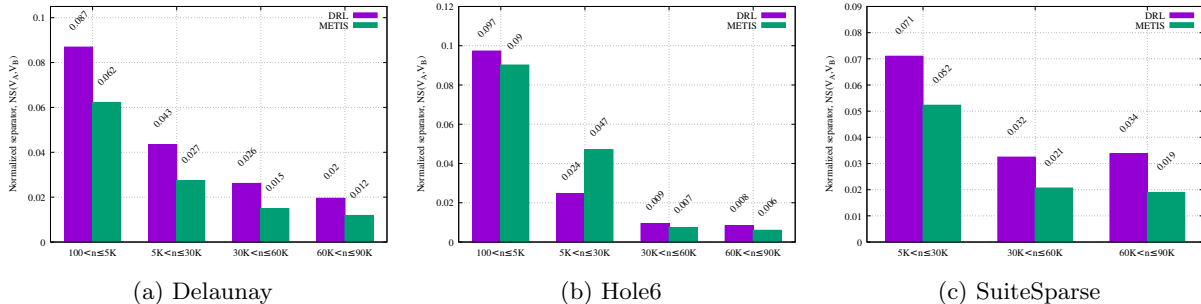


Figure 14: Comparison of the normalized separator, Eq. 10, using the proposed Algorithm 6, and METIS.




Figure 15: Nested dissection ordering of the 100×100 sparse matrix $A = \text{nos4}$ from the SuiteSparse matrix collection. It is structurally symmetric, so its sparsity pattern can be represented by a graph with 100 nodes and 594 edges. (a) The sparsity pattern of the original matrix A , with 594 number of non-zero entries (nnz). (b) The sparsity pattern of L and U , the triangular factorization of $A = LU$, with $\text{nnz} = 1417$. (c) The matrix symmetrically permuted using the minimum fill-reducing ordering: $P^T A P$, with $\text{nnz} = 594$. (d) Sparsity pattern of the LU factorization of the ordered matrix $P^T A P = LU$, with $\text{nnz} = 1160$.

separator values found with DRL are slightly higher than the ones from METIS, but overall the results are good. Indeed, the testing dataset includes larger graphs with a wide variety of sparsity patterns, many of them not included in the training dataset.

5 Nested Dissection Sparse Matrix Ordering

The nested dissection algorithm is a heuristic used to order the rows and columns of a sparse matrix before applying Gaussian elimination to the matrix. When computing an LU decomposition of a sparse matrix A as $A = LU$, the factors L (lower triangular) and U (upper triangular) are typically less sparse than the original matrix A . The extra nonzeros which are introduced in the factors are known as the fill-in. However, reordering the matrix before performing the numerical factorization can greatly reduce this fill-in, and nested dissection is known to produce orderings that drastically reduce the fill. This is illustrated in Figure 15 for the matrix `nos4`¹ from the SuiteSparse matrix collection.

Nested dissection orders the rows and columns of a sparse matrix A as follows. Consider the graph G corresponding to the sparsity pattern of A if A is symmetric, or to $A^T + A$ otherwise. Find a minimal vertex separator V_S of G , that splits G into two unconnected sets of nodes V_A and V_B . Then first order the matrix rows and columns corresponding to the nodes in V_A using nested dissection (recursively), next, number the rows and columns corresponding to the nodes in V_B using nested dissection, and lastly, order the rows and columns corresponding to the nodes in the vertex separator V_S .

Algorithm 8 shows the nested dissection algorithm, using the recursive formulation for ease of notation. However, in the actual implementation this recursion is avoided using an implementation based on a stack

¹<https://sparse.tamu.edu/HB/nos4>

Algorithm 8 Nested dissection sparse matrix reordering, using deep reinforcement learning to find a vertex separator (Algorithm 6).

Input: graph $G(V, E)$ corresponding to the sparsity pattern of a symmetric matrix A , or to $A^T + A$

Output: permutation vector p

```

1: function NESTED_DISSECTION( $G$ )
2:   if  $|V| < n_{\min}$  then
3:     return AMD( $G$ )                                ▷ end recursion by calling AMD ordering [3]
4:   end if
5:    $V_S, V_A, V_B \leftarrow \text{vertex\_separator}(G)$       ▷ See Algorithm 6
6:    $p_A \leftarrow \text{nested\_dissection}(\text{subgraph}(V_A, G))$   ▷ recursive call
7:    $p_B \leftarrow \text{nested\_dissection}(\text{subgraph}(V_B, G))$   ▷ recursive call
8:   return  $[V_A(i) \text{ for } i \text{ in } p_A] + [V_B(i) \text{ for } i \text{ in } p_B] + V_S$ 
9: end function

```

data structure. Line 6 and Line 7 show the recursive calls. The recursion is stopped early, in Line 3, by calling a different ordering algorithm once the graph becomes smaller than the threshold n_{\min} . Here we use the approximate minimum degree ordering (AMD) [3], which is a different heuristic known to produce good ordering for small to medium sized problems, and for which highly efficient sequential codes are available. Algorithm 8 returns a permutation vector p , corresponding to a permutation matrix P , which can be used to symmetrically permute the rows and the columns of the matrix.

5.1 Experimental Evaluation

We compare several sparse matrix ordering techniques on a number of sparse matrices by computing a factorization using the SuperLU sparse direct solver [36], through the Scipy Python interface [48]. The orderings considered are:

- **DRL_ND:** nested dissection using Algorithm 8, computing the vertex separators using Algorithm 6.
- **METIS_ND:** nested dissection using Algorithm 8, but computing the vertex separators using METIS.
- **METIS:** nested dissection implementation from METIS.
- **SCOTCH:** nested dissection implementation from SCOTCH.
- **COLAMD:** approximate minimum degree ordering applied to the graph of $A^T A$ (the default option for SuperLU).

In Algorithm 8, n_{\min} is set to 100.

Delaunay Triangulations Figure 16a collects the results for the fill on 100 Delaunay triangulations having between 100 and 90,000 vertices with the different orderings. Our model (DRL_ND) always outperforms minimum degree (COLAMD) and it gives results similar results to the SCOTCH orderings, while performing slightly worse than METIS and nested dissection (Algorithm 8) with METIS to find the vertex separators (METIS_ND). This shows that the model is able to generalize well to larger and unseen graphs.

Sparse Matrices from the SuiteSparse Collection Figure 16b shows the results for the fill on 101 2D/3D discretizations from the SuiteSparse matrix collection, with number of vertices between 5,000 and 90,000. We notice that our proposed algorithm (DRL_ND) performs similarly to SCOTCH and outperforms minimum degree (COLAMD) in every class of graphs. The fill with DRL turns out to be a little higher than the one obtained with METIS and METIS_ND. Recall that this testing dataset includes larger graphs than the ones used for training and a wide variety of sparsity patterns, so the results show that the model generalizes well to larger graphs, and to graphs from different types of discretizations.




Figure 16: Comparison of the fill in the triangular factors L and U of a matrix A corresponding to (Figure 16a) the adjacency matrix (and a diagonal shift) of several Delaunay graphs or (Figure 16b) several SuiteSparse matrices. For DRL_ND, the sparse matrix is reordered using Algorithm 8 with Algorithm 6 to find the vertex separators. For METIS_ND, the permutation is computed using Algorithm 8, but with METIS to find the vertex separators. The METIS datapoints in this figure use the nested dissection algorithm as implemented in the METIS package.

6 Conclusions

We have presented a graph partitioning approach based on deep reinforcement learning, using a multilevel framework. We show both an edge partitioner, computing graph bisections, as well as a variant of the method which computes a vertex separator. We show that the graph partitioning and vertex separator codes, when trained on graphs of a certain type and with less than 5,000 nodes, generalize well to graphs with many more nodes, as well as to different types of graphs. For instance, Figures 12b and 12a show that for graphs with up to 90,000 nodes, the cut and the normalized cut for the partitions from the proposed method, are very competitive with those computed using METIS and SCOTCH. Note, the graphs we used for this test are from the SuiteSparse dataset, which contains problems from a variety of applications, showing a wide range of sparsity patterns.

In Section 5, the vertex separator code is used recursively to construct a nested dissection ordering, which we then evaluate in the sparse solver SuperLU. We show that the resulting sparse matrix ordering effectively reduces the fill-in, and does this more so than the approximate minimum degree ordering. The quality of the ordering is comparable to the ordering produced by SCOTCH and only slightly worse than the nested dissection ordering from METIS. We believe further tuning of the neural network, and training on more and larger graphs, could further improve these results.

From our complexity analysis in Section 3.3, we believe that the presented approach can also perform well, competitively with the state-of-the-art codes METIS and SCOTCH. As a possible improvement to the runtime of the algorithms, we could consider selecting multiple actions at once, effectively reducing the number of forward evaluations of the neural network. Furthermore, we plan to port the method from Python to compiled C++, to gain additional computation efficacy.

Acknowledgements

This work was supported by the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under U.S. Department of Energy Contract No. DE-AC02-05CH11231.

References

- [1] Abadi, M. et al. TensorFlow: Large-scale machine learning on heterogeneous systems. <https://www.tensorflow.org/>, 2015. Software available from tensorflow.org.
- [2] Reza Refaei Afshar, Y. Zhang, Murat Firat, and U. Kaymak. A State Aggregation Approach for Solving Knapsack Problem with Deep Reinforcement Learning. In *ACML*, 2020.
- [3] P.R. Amestoy, T.A. Davis, and I.S. Duff. An Approximate Minimum Degree Ordering Algorithm. *SIMAX*, 17(4):886–905, 1996.
- [4] Robert Anderson, Julian Andrej, Andrew Barker, Jamie Bramwell, Jean-Sylvain Camier, Jakub Cervený, Veselin Dobrev, Yohann Dudouit, Aaron Fisher, Tzanio Kolev, et al. MFEM: a modular finite element methods library. arXiv:1911.09220 [cs.MS], 2019.
- [5] Bas O. Fagginger Auer and Rob H. Bisseling. A GPU algorithm for greedy graph matching. In *Facing the Multicore-Challenge II*, pages 108–119. Springer, 2012.
- [6] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Lisandro Dalcin, Alp Dener, Victor Eijkhout, William D. Gropp, Dmitry Karpeyev, Dinesh Kaushik, Matthew G. Knepley, Dave A. May, Lois Curfman McInnes, Richard Tran Mills, Todd Munson, Karl Rupp, Patrick Sanan, Barry F. Smith, Stefano Zampini, Hong Zhang, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.13, Argonne National Laboratory, 2020.
- [7] P. Battaglia, Jessica B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, A. Santoro, R. Faulkner, Çağlar Gülçehre, H. Song, A. J. Ballard, J. Gilmer, George E. Dahl, Ashish Vaswani, Kelsey R. Allen, Charlie Nash, Victoria Langston, Chris Dyer, N. Heess, Daan Wierstra, P. Kohli, M. Botvinick, Oriol Vinyals, Y. Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks. arXiv:1806.01261 [cd.LG], 2018.
- [8] Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural Combinatorial Optimization with Reinforcement Learning. arXiv:1611.09940 [cs.AI], 2017.
- [9] 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, July 2017.
- [10] Jeff Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. In *Proceedings of the Princeton conference in honor of Professor S. Bochner*, pages 195–199, 1969.
- [11] C. Chevalier and F. Pellegrini. PT-Scotch: A tool for efficient parallel graph ordering. *PARCO*, 34(6-8):318–331, 2008.
- [12] Fan R. K. Chung. *Spectral Graph Theory*. AMS, 1997.
- [13] T. Degris, P. Pilarski, and R. Sutton. Model-free reinforcement learning with continuous action in practice. *2012 American Control Conference (ACC)*, pages 2177–2182, 2012.
- [14] Karen Devine, Erik Boman, Robert Heaphy, Bruce Hendrickson, and Courtenay Vaughan. Zoltan data management service for parallel dynamic applications. *Computing in Science & Engineering*, 4(2):90–97, 2002.
- [15] Iain S. Duff. On algorithms for obtaining a maximum transversal. *ACM Transactions on Mathematical Software (TOMS)*, 7(3):315–330, 1981.
- [16] Robert D Falgout and Ulrike Meier Yang. hypre: A library of high performance preconditioners. In *International Conference on Computational Science*, pages 632–641. Springer, 2002.
- [17] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

- [18] C.M. Fiduccia and R.M. Mattheyses. A linear-time heuristic for improving network partitions. In *19th Design Automation Conference*, pages 175–181. IEEE, 1982.
- [19] Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA, 1990.
- [20] Graph Nets: DeepMind’s library for building graph networks in Tensorflow and Sonnet. https://github.com/deepmind/graph_nets. Accessed: 2020-05-20.
- [21] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pages 1024–1034, 2017.
- [22] William L. Hamilton, Rex Ying, and J. Leskovec. Representation Learning on Graphs: Methods and Applications. arXiv:1709.05584 [cs.SI], 2017.
- [23] Bruce Hendrickson and Edward Rothberg. Improving the run time and quality of nested dissection ordering. *SIAM Journal on Scientific Computing*, 20(2):468–489, 1998.
- [24] Michael A. Heroux, Roscoe A. Bartlett, Vicki E. Howle, Robert J. Hoekstra, Jonathan J. Hu, Tamara G. Kolda, Richard B. Lehoucq, Kevin R. Long, Roger P. Pawlowski, Eric T. Phipps, et al. An overview of the trilinos project. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):397–423, 2005.
- [25] Nikolaos Karalias and Andreas Loukas. Erdős goes neural: an unsupervised learning framework for combinatorial optimization on graphs. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, Vancouver, Canada, 2020.
- [26] G. Karypis and V. Kumar. Parallel Multilevel k -Way Partitioning Scheme for Irregular Graphs. *SIAM Review*, 41(2):278–300, 1999.
- [27] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.
- [28] B.W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell system technical journal*, 49(2):291–307, 1970.
- [29] Elias Boutros Khalil, H. Dai, Yuyu Zhang, B. Dilkina, and L. Song. Learning combinatorial optimization algorithms over graphs. In *NIPS*, 2017.
- [30] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.
- [31] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. arXiv:1609.02907 [cs.LG], 2017.
- [32] Andrew V. Knyazev. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM journal on scientific computing*, 23(2):517–541, 2001.
- [33] Scott P. Kolodziej, Mohsen Aznaveh, Matthew Bullock, Jarrett David, Timothy A. Davis, Matthew Henderson, Yifan Hu, and Read Sandstrom. The Suitesparse matrix collection website interface. *Journal of Open Source Software*, 4(35):1244, 2019.
- [34] Cornelius Lanczos. *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*. United States Governm. Press Office Los Angeles, CA, 1950.
- [35] Charles E. Leiserson and John G. Lewis. Orderings for parallel sparse symmetric factorization. In *Proceedings of the Third SIAM Conference on Parallel Processing for Scientific Computing*, pages 27–31, 1987.
- [36] X.S. Li. An overview of SuperLU: Algorithms, implementation, and user interface. *ACM TOMS*, 31(3):302–325, 2005.

- [37] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. arXiv:1511.05493v4 [cs.LG], 2017.
- [38] Azade Nazi, Will Hang, Anna Goldie, Sujith Ravi, and Azalia Mirhoseini. Gap: Generalizable approximate graph partitioning framework. arXiv:1903.00614 [cd.LG], 2019.
- [39] Francois Pellegrini. Distillating knowledge about SCOTCH. In Uwe Naumann, Olaf Schenk, Horst D. Simon, and Sivan Toledo, editors, *Combinatorial Scientific Computing*, number 09061 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2009. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany.
- [40] A. Pothen, H.D. Simon, and K.-P. Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIMAX*, 11(3):430–452, 1990.
- [41] Alex Pothen and Chin-Ju Fan. Computing the block triangular form of a sparse matrix. *ACM Transactions on Mathematical Software (TOMS)*, 16(4):303–324, 1990.
- [42] Marcelo O. R. Prates, Pedro H. C. Avelar, Henrique Lemos, L. Lamb, and Moshe Y. Vardi. Learning to solve NP-complete problems - a graph neural network for the decision TSP. In *AAAI*, 2019.
- [43] Horst D. Simon. Partitioning of unstructured problems for parallel processing. *Computing systems in engineering*, 2(2-3):135–148, 1991.
- [44] Barry Smith, Petter Bjorstad, and William Gropp. *Domain decomposition: parallel multilevel methods for elliptic partial differential equations*. Cambridge university press, 2004.
- [45] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.
- [46] Lloyd N. Trefethen and David Bau III. *Numerical linear algebra*, volume 50. Siam, 1997.
- [47] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [48] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [49] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992.
- [50] Alexander Zai and Brandon Brown. *Deep reinforcement learning in action*. Manning Publications, 2020.
- [51] Jiongzhi Zheng, Kun He, Jianrong Zhou, Yan Jin, and Chu-Min Li. Combining reinforcement learning with Lin-Kernighan-Helsgaun algorithm for the Traveling Salesman Problem. arXiv:2012.04461 [cs.AI], 2021.