
Optimizing Tensor Network Contraction Using Reinforcement Learning

Eli Merom¹ Haggai Maron¹ Shie Mannor¹ Gal Chechick¹

Abstract

Quantum Computing (QC) stands to revolutionize computing, but is currently still limited. To develop and test quantum algorithms today, quantum circuits are often simulated on classical computers. Simulating a complex quantum circuit requires computing the contraction of a large network of tensors. The order (path) of contraction can have a drastic effect on the computing cost, but finding an efficient order is a challenging combinatorial optimization problem. We propose a Reinforcement Learning (RL) approach combined with Graph Neural Networks (GNN) to address the contraction ordering problem. The problem is extremely challenging due to the huge search space, the heavy-tailed reward distribution, and the challenging credit assignment. We show how a carefully implemented RL-agent that uses a GNN as the basic policy construct can address these challenges and obtain significant improvements over state-of-the-art techniques in three varieties of circuits, including the largest scale networks used in contemporary QC.

1. Introduction

Quantum computing has the potential to revolutionize the computer industry (Gill et al., 2022), but quantum computers are still notoriously difficult to build (Franklin and Chong, 2004; Ladd et al., 2010; Cho, 2020). To devise and test quantum algorithms, researchers use classical computers to simulate the corresponding quantum circuits that implement them (Huang et al., 2021; Gray, 2018; Fishman et al., 2020; Patti et al., 2021b). In most cases, simulating a complex quantum circuit involves computing a contraction of a huge network of tensors.

¹Nvidia research, Israel. Code: https://nv-research-tlv.netlify.app/publication/tensor_contraction/. Correspondence to: Eli Merom <emeirom@nvidia.com>.

Proceedings of the 39th International Conference on Machine Learning, Baltimore, Maryland, USA, PMLR 162, 2022. Copyright 2022 by the author(s).

Importantly, computing the contraction can be executed in a variety of ways, all yielding the same final result, but possibly at a very different computational cost. Section 3 explains this in more detail, but for now, a good analogy is the problem of computing the product of a sequence of matrices ABC . Both $(AB)C$ and $A(BC)$ yield the same result, but the computational cost may be very different depending on the matrix dimensions.

Similarly, the cost of computing the contraction depends heavily on the order by which the tensors are contracted. Formally, the combinatorial optimization problem *Tensor Network Contraction Ordering* (TNCO) attempts to determine the optimal order of contracting a network of tensors by iteratively contracting pair by pair (Pfeifer et al., 2014; Xu et al., 2019; Schindler and Jermyn, 2020). TNCO can be naturally formulated using graphs, where nodes represent tensors and edges correspond to shared indices. Each edge is assigned a cost, which represents the computational cost of calculating the contraction along that edge. The objective is to find an ordering¹ over the pairwise contraction operations, or edges, that has the minimal total cost.

Going beyond simulations of quantum computers, tensor network contractions play a key role in other areas of science, including many-body physics (Ran et al., 2020), statistical mechanics (Liu et al., 2021), and machine learning (Stoudenmire and Schwab, 2016).

In the general case, TNCO is computationally hard (Chi-Chung et al., 1997), so designing efficient contraction algorithms remains a challenge. Some specific instances are computationally tractable, for instance, finding an optimal order for matrix multiplication, when each tensor is a matrix (Cormen et al., 2009). For more complex networks, like those relevant for quantum circuits, there are multiple approaches (Pfeifer et al., 2014; Xu et al., 2019; Schindler and Jermyn, 2020; Daniel et al., 2018), employing a variety of tools from combinatorial optimization with different quality-complexity trade-offs. The current state-of-the-art is a graph-partitioning based solver (Gray and Kourtis, 2021) that can produce high-quality solutions to networks with up to hundreds of tensors. For larger networks, the greedy approach (e.g. (Daniel et al., 2018)), which often results in

¹we use path, sequence and ordering interchangeably throughout the paper.

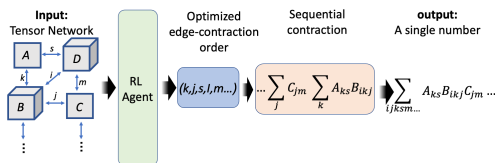


Figure 1. Contracting large tensor networks (sets of tensors with shared indices) is a major bottleneck when simulating quantum circuits. The order by which summation takes place (blue) does not affect the output, but different orders may have significantly different computation costs (orange). We tackle the Tensor Network Contraction Ordering problem: given a tensor network, our task is to find an optimal contraction order. To this end, we frame the problem as an RL problem and propose the first learning-based approach for solving it.

poor contraction paths, remains the preferred choice.

Following the recent success of ML for combinatorial optimization (Dai et al., 2017; Li et al., 2018; Sato et al., 2019; Almasan et al., 2019; Nair et al., 2020; Cappart et al., 2021; Peng et al., 2021; Meirum et al., 2021), we devise a novel and effective Reinforcement Learning (RL) approach to solve TNCO. We formulate the problem as a Markov Decision Process (MDP): The state space is defined as a space of graphs, the action space is the set of edges to be contracted, and the transition function maps a graph to a contracted graph based on the selected edge. The reward function is defined as the (negative) cost of edge contraction. To solve this problem, an RL agent starts from a graph representing the input tensor network (TN) and is trained to sequentially select edges for contraction with the goal of minimizing the cumulative contraction cost. Importantly, at each step, the agent has to balance two, seemingly conflicting, objectives: (1) Select an edge with a low cost; (2) Select an edge whose contraction yields a graph that is efficiently contractable in the future.

Notably, applying an off-the-shelf RL approach to this formulation fails due to several unique challenges: (1) The search space is huge, relative to most RL applications; (2) The reward function is strongly heavy-tailed, breaking many assumptions of classical RL algorithms; and (3) The credit assignment problem becomes very challenging due to long episodes. Finally, it is not clear how non-learning based solvers can be incorporated into the training procedure.

Aiming to address the challenges mentioned above, we introduce a novel RL approach, dubbed TNCO-RL. Our agent is implemented as a GNN (Battaglia et al., 2018) that takes as input the graph representation of the tensor network and outputs a distribution over its edges. This output distribution represents the relative likelihood that the agent selects a given edge. The agent is trained to minimize the total contraction cost using a Proximal-Policy optimization (PPO)

algorithm (Schulman et al., 2017).

Our experiments indicate that TNCO-RL outperforms state-of-the-art baselines (e.g., graph-partitioning based methods) in a variety of experimental settings on both synthetic tensor networks and tensor networks that originate from real quantum circuits (Arute et al., 2019). In summary, the contributions of this paper are: (1) We formulate the TNCO problem as a RL problem and discuss the main challenges that arise. (2) We propose a new approach to solve this problem using RL and GNNs. To the best of our knowledge, this is the first learning-based approach for the TNCO problem. (3) We apply four technical components that allow to scale the RL approach: (a) path pruning, (b) optimistic buffer, (c) baseline integration, and (d) feature robustification, (4) we provide a training environment and benchmark to allow the community to further study this problem.

2. Previous work

Tensor network contraction order optimization. Finding efficient contraction paths for tensor networks was considered in several previous works. For tiny tensor networks consisting of up to 18 tensors, one can use brute force for finding the optimal path. (Pfeifer et al., 2014) showed that optimal solutions for slightly larger networks can be obtained when using an enhanced pruning techniques. Greedy algorithms are a popular choice for larger problems (e.g., (Daniel et al., 2018)), often yielding sub-optimal paths. (Schindler and Jermyn, 2020) recently applied simulated annealing and genetic algorithms to the problem and showed that these methods outperform the greedy approach on small networks. (Gray and Kourtis, 2021) survey several additional approaches and combine them with a hyperparameter optimization approach. Moreover, they suggest using a graph partitioning algorithm (Schlag et al., 2016; Akhremtsev et al., 2017) to guide the partitioning order. Their partitioning based solver is considered the current state of the art for Sycamore networks (several hundred nodes). (Pan et al., 2021a) propose a modified Sycamore circuit TN, with a handcrafted specific ordering (rather than a general solver).

For certain specific network structures such as trees, optimal polynomial-time algorithms were suggested (Xu et al., 2019). (Huang et al., 2021) improve the actual TN contraction under memory constraint, given a contraction order. They do not discuss how to initially find a good contraction order.

RL for combinatorial optimization. Traditional algorithms often employ heuristics to solve combinatorial optimization problems. Machine learning, and RL in particular, can be used to *learn* suitable rules from a specific data distribution at hand. Over the last few years, RL approaches have been applied to multiple combinatorial optimization

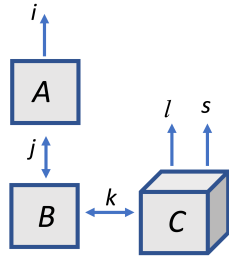
problems such as the Traveling Salesman Problem (TSP) (Bello et al., 2016), Max-Cut (Dai et al., 2017), Bin Packing Problem (BPP) (Hu et al., 2017), etc. In particular, the combination of RL techniques with agents parameterized with Graph Neural Networks (GNNs) was shown to outperform the state of the art in several cases (Dai et al., 2017; Almasan et al., 2019; Meiron et al., 2021). We refer the interested reader to two recent surveys for more information (Yang and Whinston, 2020; Mazyavkina et al., 2021).

3. Background on tensor networks

Here, we review the basic definitions of tensors and contractions. We refer the reader to (Bridgeman and Chubb, 2017) for more information about TNs and their use in QC.

Given a set of tensors, tensor network contraction refers to the operation of summing over all shared indices. Tensor networks are usually sequentially contracted by contracting a pair of tensors at a time. In that case, the cost of computing the contraction depends heavily on the order by which tensors are contracted.

Consider for example three tensors $A_{i,j}, B_{j,k}, C_{k,l,s}$ with corresponding dimensions I, J, K, L (see inset). Their contraction is defined as a fourth tensor $T_{i,l,s}$ defined by:



$$T_{i,l,s} = \sum_{j,k} A_{i,j} B_{j,k} C_{k,l,s} \quad (1)$$

If we first contract j and then k , $IJK + IKLS$ flops are required to execute the contraction. However, if we first contract k followed by j , the cost is $JKLS + IJLS$. For instance, if $I, L, S = 2, J = 10$ contracting the network with the first order costs 56 flops, while contracting with the second order costs 160 flops. With large tensor networks, the cost of different contraction orders (“paths”), can vary by orders of magnitude.

Tensors. A tensor A is a multidimensional array of numbers indexed by a subset of a finite symbolic index set $\mathcal{I}_A \subseteq \mathcal{I} = \{i, j, \dots\}$. For each index $i \in \mathcal{I}$ we define a valid range, i.e. $i \in \{1, \dots, I\}$ called the *extent* of the index.

Pairwise Tensor contraction is a common binary operation defined on pairs of tensors. Formally, the contraction of tensors A, B is a new tensor \mathcal{C} , indexed by the symmetric difference of the index sets $\mathcal{I}_A \Delta \mathcal{I}_B := (\mathcal{I}_A \cup \mathcal{I}_B) \setminus (\mathcal{I}_A \cap \mathcal{I}_B)$ whose values are defined by the following summation

on the intersection of the index sets:

$$\mathcal{C}_{\mathcal{I}_A \Delta \mathcal{I}_B} = \sum_{\mathcal{I}_A \cap \mathcal{I}_B} A_{\mathcal{I}_A} \cdot B_{\mathcal{I}_B} \quad (2)$$

Here, with a slight abuse of notation, $A_{\mathcal{I}}$ denotes the indexing of a tensor by the indices in \mathcal{I} , and $\sum_{\mathcal{I}}$ represents summation over all the indices in \mathcal{I} (see the matrix product example below). In words, the tensor contraction operation is indexed by the indices that appear in only one of the tensors and its values are a summation over all shared indices. Matrix multiplication is a simple example of tensor contraction. In that case, assuming that a matrix A is indexed by $\mathcal{I}_A = \{i, k\}$ and a matrix B is indexed by $\mathcal{I}_B = \{k, j\}$ then the contraction $\mathcal{C}(A, B)$ is a tensor indexed by $\mathcal{I}_A \Delta \mathcal{I}_B = \{i, j\}$ and defined as:

$$\mathcal{C}(A, B)_{i,j} = \sum_{k=1}^K A_{i,k} B_{k,j},$$

which coincides with matrix multiplication.

Tensor networks. A tensor network is a representation of a tensor contraction problem. Formally, a tensor network is a collection of n tensors $\{A^k\}_{k=1}^n$ where each tensor is indexed by subsets of an index set $\mathcal{I}_k \subseteq \mathcal{I}$.

It is convenient to define a graph structure $G = (V, E, w)$ on a tensor network. This is done by first defining the set of nodes V to be the set of tensors $\{A^k\}_{k=1}^n$. We then define the connectivity of the graph according to the shared index structure, namely $e_{kk'} \in E$ if $\mathcal{I}_k \cap \mathcal{I}_{k'} \neq \emptyset$. Lastly, we define a weight function on the edges $w : E \rightarrow \mathbb{R}_+$. The weight function represents some measure of the time or space complexity of contracting the indices associated with the edge, and is usually a function of the relevant extents. In this paper, w represents the number of flops needed to execute a contraction.²

Tensor network contraction. The contraction of a tensor network is a direct generalization of the binary contraction defined in Equation 2 above to a collection of tensors, namely:

$$\mathcal{C}(A^1, \dots, A^n)_{\mathcal{I}_{\text{distinct}}} = \sum_{\mathcal{I}_{\text{shared}}} A^1_{\mathcal{I}_1} \cdots A^n_{\mathcal{I}_n} \quad (3)$$

Here, $\mathcal{I}_{\text{shared}}$ is the subset of \mathcal{I} which holds indices that appear twice in the collection A^1, \dots, A^n , and $\mathcal{I}_{\text{distinct}} = \cup \mathcal{I}_i \setminus \mathcal{I}_{\text{shared}}$ holds indices that appear once.³

²The number of flops can be calculated easily as it is proportional to product of the extents of the indices that appear in both tensors.

³We assume that the total number of on index appearances is bounded by 2.

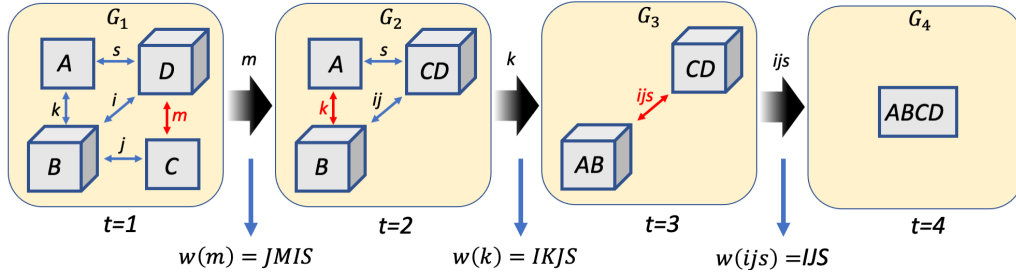


Figure 2. A sequence of pairwise contractions leading to a full contraction of the initial tensor network $G = G_1$. At each stage the red edge is contracted. w represents the cost of each contraction step.

Executing the contraction in Equation 3 directly by using nested loops over all indices is severely inefficient due to the exponential complexity of that procedure. A more efficient option is to execute the contraction as a sequence of edge (or pairwise) contractions, $P = (e_1, \dots, e_{n-1})$.⁴ Here, each contraction yields a new graph G_t , starting with $G_1 = G$. At each step t , a pair of tensors connected by the edge $e_t = (u, v)$ is contracted according to Equation (2). Then a new tensor network $G_{t+1} = (V_{t+1}, E_{t+1}, w_{t+1})$ is defined by removing the two contracted tensors u, v and adding the resulting tensor T_{t+1} to V_{t+1} . The adjacency structure is defined according to shared indices as described above (i.e., all edges that have previously been connected to either u or v will now be connected to T_{t+1} instead), see Figure 2.

4. TNCO: A definition and a RL formulation

We now define the TNCO problem and formulate it as an RL problem. We first motivate our approach, then formulate TNCO as an MDP, and discuss two possible learning setups that are of interest.

4.1. Motivation

Recently, reinforcement learning has been shown to be a promising approach for solving combinatorial optimization problems (Yang and Whinston, 2020; Mazyavkina et al., 2021). The main idea is to break down the assignment of variables in the problem to a sequential process and formulate it as a Markov Decision Process. If the problem instance at hand is similar to other instances, a learning approach may succeed in generalizing past experience and problems and achieving a good solution.

Many combinatorial optimization problems can be formulated using a graph structure. In particular, TNCO can always be written as a graph under the common assumptions in the literature (Gray and Kourtis, 2021). The most successful current approaches for TNCO use graph partitioning

⁴this approach may use more memory for storing intermediate results, but often has better time complexity.

tools (Gray and Kourtis, 2021). Graph Neural Networks (GNNs) have been shown to learn useful representations of graphs (Zhou et al., 2020). Therefore, it makes sense to combine their representative power with an RL framework.

As an important alternative, Genetic Algorithms (GAs) have been considered for the TNCO problem (e.g. (Schindler and Jermyn, 2020)). Unfortunately, they suffer from two main drawbacks. First, the search space is huge, requiring very large generations, and that makes GAs very slow and limits them to TNs of a few tens of nodes. Second, GAs do not encode the structure of the graph, so the similarity between graphs from the same distribution is not used.

4.2. The TNCO problem

Given a tensor network represented as a graph $G = (V, E, w)$, a contraction path $P = (e_1, \dots, e_{n-1})$, $e_t \in E_t$ and a corresponding sequence of graphs (G_1, \dots, G_n) , we define the total contraction cost c as the sum of weights of the contracted sequence:

$$c(P) = \sum_{t=1}^{n-1} w_t(e_t) \quad (4)$$

The goal in TNCO is to find a sequence of edge contractions P of minimal cost, namely:

$$P^*(G) = \operatorname{argmin}_P c(P) \quad (5a)$$

$$\text{s.t. } P = (e_1, \dots, e_{n-1}), e_t \in E_t \quad (5b)$$

4.3. MDP Formulation

To formulate TNCO as an MDP, we specify the state and action spaces, as well as the transition dynamics and the cost function. We use the cost rather than the reward to adhere to the QC community standards. To define an MDP, we need to define the state space \mathcal{S} ; the action space \mathcal{A} ; the transition function \mathcal{F} ; and the immediate reward/cost function \mathcal{R} . We discuss the design choice of the various elements below.

The state space \mathcal{S} is the space of all weighted graphs, $\mathcal{S} = \{G = (V, E, w)\}$. This is a very large space, but

it provides a sufficient statistic to evaluate the dynamics of every action. Also, by choosing such a state space, we open the possibility for generalization over states, that is, generalization over graphs. The actions available in a given state is the set of edges E , which naturally varies from one graph (state) to another. The transitions in this problem are deterministic: $\mathcal{F}(G, e)$ takes as input the current state G and an edge $e \in E$, and outputs a new graph that represents the tensor network after contracting the edge e as explained in Section 3. The immediate cost $\mathcal{R}(G, e)$ for an action e is the weight of the edge $w(e)$. The agent’s objective is to find the best contraction sequence, that is, to minimize the overall cost.

4.4. The learning setup: Single and multi- network learning

We consider two learning setups. First, training an agent on a single network. This setup is suitable for optimizing large networks, which is typically done once, as in quantum-supremacy tests. The agent would repeatedly contract the same network and learn to find better contraction paths. The second setup involves training the agent on a set of training networks and then applying it at inference time to new networks that are somewhat similar to the training networks, for example, taken from the same distribution.

Problem 1: Learning in a single network. Given a tensor network represented as a graph G we look for an agent H in the class of agents \mathcal{H} whose edge contraction choices yield an optimal contraction sequence. Formally, the agent should solve the following optimization problem:

$$\operatorname{argmin}_{H \in \mathcal{H}} c(P) \quad (6a)$$

$$\text{s.t. } P = (e_1, \dots, e_{n-1}) \quad (6b)$$

$$e_t = H(G_t), \quad t = 1, \dots, n-1 \quad (6c)$$

Figure 3 illustrates the edge selection process.

Problem 2: Learning and generalization to new networks Here, an agent is trained on multiple input graphs $\{G^i\}_{i=1}^m$ sampled from an underlying tensor network distribution. The agent is tested on unseen target graphs from the same distribution. In this case, the training objective is:

$$\operatorname{argmin}_{H \in \mathcal{H}} \sum_{i=1}^m c(P^i) \quad (7a)$$

$$\text{s.t. } P^i = (e_1^i, \dots, e_{n-1}^i) \quad (7b)$$

$$e_t^i = H(G_t^i), \quad t = 1, \dots, n^i - 1 \quad (7c)$$

This setup is suitable if one is interested in quickly solving TNCO instances from the same distribution. Once a model has been trained, it can be applied very quickly to infer paths for new input tensor networks.

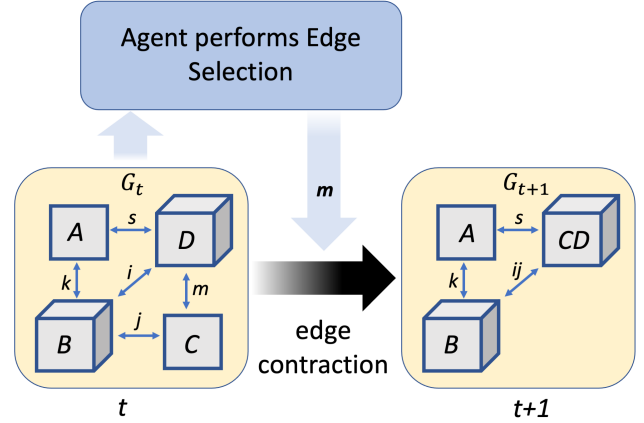


Figure 3. Our goal is to find an agent H that takes the current graph G_t as input and selects the next edge to contract, such that the sequence of selected edges P has optimal cost $c(P)$.

5. Challenges in Applying RL to TNCO

Applying RL to TNCO raises five challenges that need to be addressed. Section 6 describes our concrete solutions to these challenges in the context of TNCO.

Challenge 1: Wide dynamic range.

The contraction costs and tensor sizes typically vary across many orders of magnitude.

Deep neural networks are hard to train when costs are highly variable.

Importantly, naive

normalization does not address the problem. When the distribution of costs remains strongly unbalanced, normalizing by any constant quantity, like the max or median, washes out the signal in (at least) one end of the feature-value distribution. The inset Fig. 4 shows a histogram of edge contraction costs (in log10 scale) sampled from random paths on random TNs with $n \in \{25, 50, 75, 100\}$ tensors.

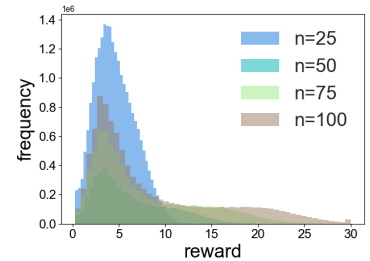


Figure 4. Reward distribution of random actions in random TNs.

Challenge 2: Huge search space.

The number of possible contraction paths in our problem is enormous. Even among very large action and state spaces, the TNCO problem is an outlier. For comparison, the average length of an episode in a game of GO (Silver et al., 2017) is 200, and so is the action space (branching factor). In a TNCO problem based on a modern quantum circuit, the action space and the length of

an episode can reach $O(1000)$ steps and actions. Therefore, the number of possible trajectories in TNCO is *hundreds of orders of magnitude larger* than that of GO.

Challenge 3: Slow convergence due to Heavy-tailed cost distribution. Nearly all research in RL is based on the assumption that the immediate cost has a sub-Gaussian distribution. For approaches based on stochastic approximation, like Q-learning or policy gradient, a bounded second moment is needed. In TNCO, as shown above for Challenge 1, the reward seems to have a heavy-tailed behavior, with its empirical distribution resembling a log-normal distribution. This implies that the variance is expected to be very large even if finite. Moreover, this phenomenon becomes worse as the networks become larger. We note that previous results for multiarmed bandits and RL (Bubeck et al., 2013; Zhuang and Sui, 2021) have shown that if the reward distribution is not sub-Gaussian, the convergence time is $|\mathcal{S}|^\alpha$, $\alpha > 1$, where α is a function of the highest moment of the distribution that converges. Consequently, the TNCO problem suffers from a huge state space and a heavy-tailed reward distribution, and therefore vanilla policy gradient algorithms are unlikely to converge in a reasonable time.

Challenge 4: Incorporating existing solvers. Existing meta-heuristic approaches and learning-based approaches have complementary strengths. Existing solvers can often find the optimal solution for small problems. Learning-based methods usually only approximate the optimal solution, but work well on large problems. The challenge is to enjoy the benefits of the two approaches.

Challenge 5: Credit assignment problem. In TNCO, as in other combinatorial optimization problems, early decisions can have a crucial effect on the overall trajectory cost. TNCO’s lengthy episodes compared to many other RL problems exacerbate the credit assignment problem.

6. Our Approach: RL-TNCO

We start with a short overview, and then discuss the graph representation and processing module in our approach. Finally, we review the RL framework and focus on how we address the challenges discussed in the previous section.

Overview. As the basis for our RL framework, we use PPO (Schulman et al., 2017), though the choice of the specific policy search algorithm is less important here. Tensor networks are represented as graphs, and we use a GNN to parameterize our policy (Battaglia et al., 2018). At each step, a graph representation of the current tensor network G is fed into the GNN, and the GNN outputs a distribution over the edges to be selected. An edge is then sampled from the distribution and contracted. This yields a smaller tensor

network with one less node, represented by the graph G' . The weight of the selected edge is added to the total path cost, which is used to calculate the loss. The process is repeated until no edges remain in G' . At inference time, we treat the agent output as a distribution over contraction sequences. We sample ℓ sequences, edge by edge, and use the sequence that yielded the minimal total contraction cost.

6.1. Graph representation and architecture

We represent the graph as a tuple (A, X, E, g) where $A \in \{0, 1\}^{n \times n}$ is the adjacency matrix of the tensor network, $X \in \mathbb{R}^{n \times d}$ is a node feature matrix that holds features for each tensor, $E \in \mathbb{R}^{m \times d'}$ is an edge feature matrix, and $g \in \mathbb{R}^{d''}$ is a global graph feature vector. Here, m is the number of edges in the tensor network and d, d', d'' are arbitrary feature dimensions. We experimented with several choices of initial features and converged on the following choice: For node and global features we use constant input features, i.e. $X = [1, \dots, 1]$, $g = [1]$. The edge features of an edge $e = (u, v)$ are: (1) its are the contraction cost; (2) the edge’s extension (product of the extents of the associated indices), and (3) the greedy score of the edge $e = (u, v)$, based on a heuristic approach in OPT-Einsum (Daniel et al., 2018) $\prod_{i \in \mathcal{I}_u} I(i) + \prod_{i \in \mathcal{I}_v} I(i) - \prod_{i \in \mathcal{I}_u \Delta \mathcal{I}_v} I(i)$. Here, $I(i)$ is the extent of index i .

As we have both node and edge features, we choose to use a popular message passing GNN model (Battaglia et al., 2018) (Algorithm 1, pp. 12) that maintains learnable node, edge and global graph representations. The model is composed of several message passing layers, $L^k \circ \dots \circ L^1$ where each L^i updates all the representations, i.e.:

$$X^{i+1}, E^{i+1}, g^{i+1} = L^i(A, X^i, E^i, g^i; \theta^i),$$

Each layer L_i updates the features sequentially: node and edge features are updated by aggregating local information, while the global feature is updated by aggregating over the whole graph. We denote the parameters of the MLPs that are used in a layer L_i as θ_i , and note that these are the only learnable parameters in the model. At the last layer $i = k$ we use a single dimension for edge features, that is, $d' = 1$ and apply a normalization function over edge scores $s(e)$ (Meirom et al., 2021),

$$\Pr(e) = \frac{s'(e)}{\sum_e s'(e)}, \quad s'(e) = s(e) - \min_{e'} s(e') + \epsilon.$$

This provides us with distribution $\Pr(\cdot)$ over the edges, which is used for sampling an edge. The critic module shares the same architecture, on top of which a final pooling layer over the nodes is added to obtain the value function.

6.2. Algorithmic modifications

To address the challenges discussed in Section 5, we incorporated several algorithmic modifications to vanilla deep policy gradient algorithms.

Path pruning. To restrict the huge search space (Challenge #2), we use path pruning. All the costs in a TNCO problem are strictly positive. Therefore, if the cost accumulated at some point in a path exceeds the smallest observed cost C^{min} , this path cannot be optimal and can be pruned.

Unlike inference, using path pruning in training cannot be done in a naive way. First, if we prune paths during the path-collection phase, the vast majority of paths get pruned and too little data is left for training. A better approach would be to use the path up to the termination point. In this approach, we need to estimate the cost of that path suffix, and this requires estimating the future expected return at the termination state, which is commonly done using the value function. However, estimating that value itself relies on the assumption that we can collect enough full (completed) trajectories, so the critic can learn to approximate the value at the tail of the trajectory. Since in TNCO most trajectories terminate prematurely, the approximation of the value function at the termination state becomes very poor. To mitigate these issues, we instead approximate the termination value after l contractions by $\min C(G, e) + (\frac{C_l}{l} + \frac{C^{min}}{n}) \cdot \frac{n-l}{2}$. The last term averages the mean known cost per step and the mean cost per step of the current trajectory $\frac{C_l}{l}$.

Optimistic buffer. In our setup, good paths are very rare. Even when the agent encounters a good path, it is often overwhelmed by poor trajectories, and the information is often washed out by negative path updates, leading to slow convergence (Challenge #3). Furthermore, deep policy gradient algorithms often apply gradient clipping, limiting the effect of the highly valuable but very rare trajectories. To overcome this and address the slow convergence rate of the PG algorithm in our setup, we add an off-policy element to the on-policy PPO algorithm – an optimistic buffer.

This buffer contains tuples $(G, w(e), e, C(G, e))$ ⁵, where $C(G, e)$ is the total empirical accumulated cost from this state after performing the action e , as logged from a specific trajectory. This is a pessimistic estimation of the optimal value function $V^{opt}(G)$, as the optimal value function is the maximal accumulated reward over all trajectories from G . After each training step, we update this buffer. For each sample such that $V(G) > C(G, e)$ from either the roll-out buffer or the optimistic buffer we assign a score $V(G) - C(G, e)$, while for all other samples we assign a score of 0, effectively eliminating them. We then sample b samples according to the normalized score distribution,

⁵In standard RL notation, $G = s, w(e) = r, e = a$.

where b is the optimistic buffer size. Then, we update the actor and critic modules using PPO loss with samples drawn from this buffer. This results in an off-policy bias, which shifts both the actor (and critic) towards the optimal policy.

Robust feature scaling. To address the problem of a wide dynamic range (Challenge #1), we propose a novel scheme of feature robustification. We apply two transformations to each input feature that assign high dynamic range to the tails. In addition, we apply a logarithmic transformation that compresses the dynamic range of the heavy-tail distribution. These transformations are applied to every feature. For a specific feature z , denote the vector z values across all nodes as y_z . This is a column vector of the edge feature matrix E . If all the values in y_z have the same sign, the scale feature is $\log(abs(y_z))$, otherwise it is $\log(y_z - \min(y_z) + 1)$. Next, we apply the two transformations that enhance the dynamic range at the distribution tails. We use the median m of y_z as a “border” to separate high and low ranges. We scale each column as $\frac{y_z - m}{h(y_z) - m}$, where $h(\cdot)$ is the scale of the high (low) ranges, e.g. $h(\cdot) = \max(\cdot), \min(\cdot)$. Then, we clip the high and low regimes to $[-100, 100]$.

Leveraging existing solvers. To benefit from existing algorithms (Challenge #4), which can often find optimal solutions in small networks, we tested two approaches to incorporate existing solvers. (1) For a solver based on action scores (such as a greedy approach), we incorporated its score as additional features that the algorithm can learn to use. (2) Use a fast solver to optimize part of a problem. For example, when solving a large TNCO problem, use a greedy approach to find a prefix of the contraction path, using the low-cost contractions, and then use RL to find the rest. In the later approach, RL spends optimization cycles on the most important (heavy) late stage contractions.

The intuition behind using greedy at the prefix is that the mean cost of contraction grows during optimization, hence most headroom to improve is at the suffix. For example, Fig. 5 presents the cost per step of the greedy solver for the max-cut TN.

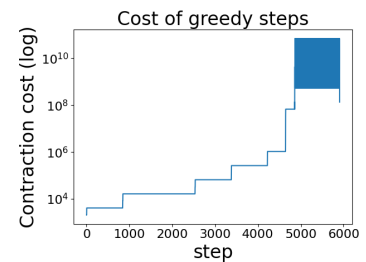


Figure 5. The contraction cost per step of the greedy algorithm on the max-cut TN. 90% of the cost incurred during the last 100 steps, while 95.5% of the contraction cost is incurred during the last 1000 steps.

We also implemented applying the solver at the suffix, but since the prefix graph is very large (10K edges), only small

batches fit in memory even on a 40GB GPU. As a result, applying RL-TNCO to the prefix is, unfortunately, slower by orders of magnitude.

The credit assignment problem (Challenge 5) is addressed through the use of a general-purpose RL scheme and a lot of compute power. Developing an algorithmic approach to address this challenge is an open problem.

7. Experiments

7.1. Setup

We conducted experiments on three types of networks: synthetic networks (≤ 100 tensors), Sycamore circuit networks (100 – 400) and Max-cut (> 5000) networks. The latter two networks originate from real quantum circuits.

When multiple tensor networks drawn from a shared distribution are available, it is beneficial to train a model to solve circuits from that distribution. Synthetic networks allowed us to test how well our method generalizes to new networks from the same distribution. Therefore, we applied the multi-network setup on the synthetic networks.

The Sycamore and Max-cut are on the frontline of the quantum supremacy regime, and are simulated in supercomputers. The key challenge here is to reduce the total flop count, whereas the runtime of the TNCO solver is of less concern. The number of both Sycamore circuits and max-cut graphs is very low (four and one, respectively), and therefore we used the single-network approach in these TNs.

Baselines. We compare our approach to three baselines: (1) **OE-GREEDY**: the greedy solver from OPT-Einsum (Daniel et al., 2018). This solver provides a very fast solution which is usually sub-optimal. It runs for a few milliseconds for small networks, a few seconds for huge ones; (2) **CTG-Greedy**: the greedy solver from (Gray and Kourtis, 2021); and (3) **CTG-Kahypar**, a strong graph partitioning based solver from (Gray and Kourtis, 2021). Both CTG-Greedy and CTG-Kahypar run with a time limit as specified in each experiment. Since OE-Greedy is very efficient, if another algorithm has a worse solution, we replace it with the OE-Greedy solution.

Implementation and system. We used NVIDIA DGX-V100 for all experiments. We used the pyTorch Geometric (Fey and Lenssen, 2019) for GNN implementations and Stable Baselines (Raffin et al., 2019) for the RL framework.

Hyper-parameter search. We conducted a hyper-parameter search on the following choices: (1) Number of GNN layers $k \in \{3, 4, 6\}$; (2) Using PairNorm or not (Zhao and Akoglu, 2019); (3) Neighborhood aggregation type (max/mean/sum); (4) Edge normalization method (soft-

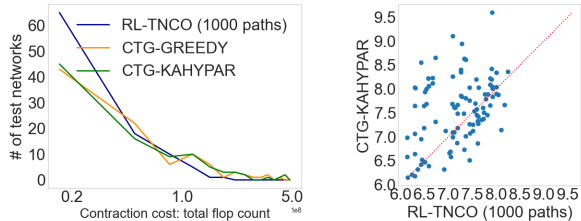


Figure 6. Contraction costs for synthetic networks of $n = 50$, scale is 10^8 . (left) A histogram of the total flop count. (right) Each point corresponds to a network from the test set, axes correspond to contraction costs for RL-TNCO and Kahypar.

max vs. the approach from (Meirom et al., 2021)).

7.2. Results

Synthetic networks. We first applied our approach to synthetic networks with sizes $n = 25, 50, 75$ and 100 nodes. Following (Gray and Kourtis, 2021) we generate graphs using the OPT-Einsum package (Daniel et al., 2018) with an average degree $d = 3$ and tensor extents that are sampled i.i.d. from a uniform distribution on $\{2, 3, 4, 5, 6\}$. In all experiments, we train on randomly generated tensor networks and test on a set consisting of specific 100 equations. For each equation, we sampled numerous paths, and report the best found path. We varied (and report) the number of paths sampled at test time $\ell \in \{50, 1000\}$. This allows us to trade off runtime with solution quality. We repeat each experiment with three random restarts (seeds) and report the average solution. CTG-Kahypar and CTG-Greedy ran for 300 seconds for each TN.

The results are reported in Table 1. RL-TNCO outperforms all baselines by a large margin ($2\times-5\times$), even when sampling only $\ell = 50$ sampled contraction sequences. Figure 6 (top) provides a more detailed analysis of the results showing that the RL-TNCO algorithm tends to provide more efficient solutions. Figure 6 (bottom) shows a paired comparison of the results of RL-TNCO and CTG-Kahypar. RL-TNCO is better than CTG-Kahypar in 72% of the networks (points below the line). Table 2 compares the running time of all algorithms at inference time. RL-TNCO runs significantly faster than the CTG baselines.

Finally, we tested how RL-TNCO models trained on tensor networks of one size perform when tested with tensor networks of a different size. This *size-generalization* problem is an active research field (Yehudai et al., 2021) We trained four models on networks with different fixed sizes, having 25, 50, 75, and 100 nodes. We tested how each model performs on a test set containing TNs with 25 nodes. We find that models trained on TNs with 50, 75, and 100 nodes reached 99.7%, 94.1%, 91.8% of the performance of “same-size” models, trained on TNs of size 25. This shows that

Table 1. Total flop count in synthetic networks of various sizes.

n	Mean \pm std				Median \pm MAD			
	25	50	75	100	25	50	75	100
Scale	10^4	10^7	10^{10}	10^{12}	10^4	10^7	10^{10}	10^{12}
OE greedy	53.7 ± 83.7	75.4 ± 253	104 ± 410	5296 ± 38310	27.7 ± 15.6	11.3 ± 10.0	4.5 ± 4.4	26.4 ± 25.4
CTG greedy	40.3 ± 70.9	12.8 ± 45.6	8.3 ± 27.1	27.9 ± 103.0	20.3 ± 9.6	4.2 ± 3.4	0.9 ± 0.8	2.2 ± 2.0
CTG kahypar	46.4 ± 76.5	13.4 ± 42.2	4.1 ± 18.2	54.2 ± 463	24.8 ± 12.4	4.3 ± 3.5	0.4 ± 0.4	1.2 ± 1.1
RL-TNCO (50)	13.9 ± 7.3	3.9 ± 4.5	1.6 ± 4.6	9.1 ± 21.1	13.5 ± 7.3	2.0 ± 4.5	0.2 ± 4.6	2.3 ± 21.1
RL-TNCO (1000)	13.1 ± 6.5	3.2 ± 3.5	1.2 ± 2.7	5.5 ± 11.0	12.5 ± 3.2	1.8 ± 1.5	0.2 ± 0.1	1.8 ± 1.6

Table 2. Inference time comparison in seconds for synthetic networks.

# nodes	25	50	75	100
OE greedy	0.02 s	0.07 s	0.1 s	0.06 s
CTG greedy	300 s	300 s	300 s	300 s
CTG kahypar	300 s	300 s	300 s	300 s
RL-TNCO (50)	1.6 s	4.2 s	6.8 s	9.6 s
RL-TNCO (1000)	33 s	74.6 s	124.2 s	162.4 s

 Table 3. Contraction cost for Sycamore networks of four sizes, with $m = 10, 12, 14, 20$. Values indicate flops count.

	m=10	m=12	m=14	m=20
# Tensors	162	210	244	379
Scale	10^{10}	10^{12}	10^{14}	10^{18}
OE greedy	$5.71 \cdot 10^4$	$2.96 \cdot 10^8$	$1.52 \cdot 10^4$	$2.04 \cdot 10^{13}$
CTG greedy	3.78	102	19.2	8.57
CTG kahypar	2.01	43.6	5.07	5.83
RL-TNCO	5.44	7.4	2.63	3.50

performance deteriorates slightly when the train distribution deviates from the test distribution.

Sycamore networks. We tested our framework on four circuits based on the sycamore processor (Arute et al., 2019), with 53 qubits and $m = 10, 12, 14, 20$ cycles. These networks had 162-379 tensors after simplification using the QUIMB library (Gray, 2018).⁶ We consider the single-circuit learning setup. We ran CTG-Greedy and CTG-kahypar for 3 hours on each circuit. RL-TNCO outperforms the baselines in three out of four TNs. See Table 3 for a comparison of our results to the baselines, and Table ?? for a quantification of gains for Sycamore $m = 20$.

Max-cut networks. We applied RL-TNCO to a tensor network that originated from a Max-Cut problem (Patti

⁶With this simplifications, the circuits are not identical to those reported in (Gray and Kourtis, 2021).

Table 4. Comparison of cost and runtime for Max-cut TNs.

Method	contraction cost $\times 10^{12}$ flops	time
OE-Greedy	36.45	5.1 sec
CTG-Kahypar	144	6 hours
CTG-Greedy	2540	6 hours
RL-TNCO (last 100 steps)	33.40	20 min
RL-TNCO (last 1000 steps)	19.05	1.5 hrs
RL-TNCO (last 2000 steps)	9.29	5.5 hrs

et al., 2021a). The circuit consists of 1688 qubits with depth 21, yielding a network with 5908 tensors. Due to the large network size and as discussed in Section 6, we first execute a greedy solver and use RL-TNCO for optimizing the last k contraction steps. Table 4 summarizes the results. Using RL-TNCO to optimize the last $k = 100, 1000$, or 2000 steps after a greedy path prefix reduces the total flop count by factors of $\times 1.09$, $\times 1.9$, and $\times 3.9$, respectively. CTG-greedy and CTG-kahypar ran with a limit of 6 hours.

8. Conclusions

In this paper, we introduce the TNCO problem to the machine-learning community and propose the first ML-based method for solving it. We formulated the TNCO problem as an MDP, discussed the main challenges that arise from this formulation, and finally devised a novel approach based on RL and GNNs. The method outperforms the state of the art on several tensor network types.

(Huang et al., 2021) showed that performing a quantum simulation of recent large scale circuits, including Sycamore circuits, requires roughly 20 days on a cluster whose size is comparable to Summit. We show here that using RL to optimize contraction order can further reduce computing time by a factor of 2x-4x, minimizing highly expensive and scarce super-computing resources, reducing the amount of energy consumption of simulating just a single circuit by GWhs, amounting to powering hundreds of household per year, saving tens of thousands of dollars, and consequently accelerating and advancing research in QC.

References

- Akhremtsev, Y., Heuer, T., Sanders, P., and Schlag, S. (2017). Engineering a direct k-way hypergraph partitioning algorithm. In *2017 Proceedings of the Nineteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 28–42. SIAM.
- Almasan, P., Suárez-Varela, J., Badia-Sampera, A., Rusek, K., Barlet-Ros, P., and Cabellos-Aparicio, A. (2019). Deep reinforcement learning meets graph neural networks: Exploring a routing optimization use case. *arXiv preprint arXiv:1910.07421*.
- Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G., Buell, D. A., et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al. (2018). Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Bello, I., Pham, H., Le, Q. V., Norouzi, M., and Bengio, S. (2016). Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*.
- Bridgeman, J. C. and Chubb, C. T. (2017). Hand-waving and interpretive dance: an introductory course on tensor networks. *Journal of physics A: Mathematical and theoretical*, 50(22):223001.
- Bubeck, S., Cesa-Bianchi, N., and Lugosi, G. (2013). Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717.
- Cappart, Q., Chételat, D., Khalil, E., Lodi, A., Morris, C., and Veličković, P. (2021). Combinatorial optimization and reasoning with graph neural networks. *arXiv preprint arXiv:2102.09544*.
- Chi-Chung, L., Sadayappan, P., and Wenger, R. (1997). On optimizing a class of multi-dimensional loops with reduction for parallel execution. *Parallel Processing Letters*, 7(02):157–168.
- Cho, A. (2020). The biggest flipping challenge in quantum computing. *Science*, 10.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2009). *Introduction to algorithms*. MIT press.
- Dai, H., Khalil, E. B., Zhang, Y., Dilkina, B., and Song, L. (2017). Learning combinatorial optimization algorithms over graphs. *arXiv preprint arXiv:1704.01665*.
- Daniel, G., Gray, J., et al. (2018). Opt_einsum—a python package for optimizing contraction order for einsum-like expressions. *Journal of Open Source Software*, 3(26):753.
- Fey, M. and Lenssen, J. E. (2019). Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*.
- Fishman, M., White, S. R., and Stoudenmire, E. M. (2020). The ITensor software library for tensor network calculations.
- Franklin, D. and Chong, F. T. (2004). Challenges in reliable quantum computing. In *Nano, quantum and molecular computing*, pages 247–266. Springer.
- Gill, S. S., Kumar, A., Singh, H., Singh, M., Kaur, K., Usman, M., and Buyya, R. (2022). Quantum computing: A taxonomy, systematic review and future directions. *Software: Practice and Experience*, 52(1):66–114.
- Gray, J. (2018). quimb: A python package for quantum information and many-body calculations. *Journal of Open Source Software*, 3(29):819.
- Gray, J. and Kourtis, S. (2021). Hyper-optimized tensor network contraction. *Quantum*, 5:410.
- Hu, H., Zhang, X., Yan, X., Wang, L., and Xu, Y. (2017). Solving a new 3d bin packing problem with deep reinforcement learning method. *arXiv preprint arXiv:1708.05930*.
- Huang, C., Zhang, F., Newman, M., Ni, X., Ding, D., Cai, J., Gao, X., Wang, T., Wu, F., Zhang, G., et al. (2021). Efficient parallelization of tensor network contraction for simulating quantum computation. *Nature Computational Science*, 1(9):578–587.
- Ladd, T. D., Jelezko, F., Laflamme, R., Nakamura, Y., Monroe, C., and O’Brien, J. L. (2010). Quantum computers. *nature*, 464(7285):45–53.
- Li, Z., Chen, Q., and Koltun, V. (2018). Combinatorial optimization with graph convolutional networks and guided tree search. *arXiv preprint arXiv:1810.10659*.
- Liu, J.-G., Wang, L., and Zhang, P. (2021). Tropical tensor network for ground states of spin glasses. *Physical Review Letters*, 126(9):090506.
- Mazyavkina, N., Sviridov, S., Ivanov, S., and Burnaev, E. (2021). Reinforcement learning for combinatorial optimization: A survey. *Computers & Operations Research*, page 105400.
- Meirom, E., Maron, H., Mannor, S., and Chechik, G. (2021). Controlling graph dynamics with reinforcement learning and graph neural networks. In *International Conference on Machine Learning*, pages 7565–7577. PMLR.

- Nair, V., Bartunov, S., Gimeno, F., von Glehn, I., Lichocki, P., Lobov, I., O’Donoghue, B., Sonnerat, N., Tjandraatmadja, C., Wang, P., et al. (2020). Solving mixed integer programs using neural networks. *arXiv preprint arXiv:2012.13349*.
- Pan, F., Chen, K., and Zhang, P. (2021a). Solving the sampling problem of the sycamore quantum supremacy circuits.
- Pan, F., Chen, K., and Zhang, P. (2021b). Solving the sampling problem of the sycamore quantum supremacy circuits. *arXiv preprint arXiv:2111.03011*.
- Patti, T. L., Kossaifi, J., Anandkumar, A., and Yelin, S. F. (2021a). Variational quantum optimization with multi-basis encodings. *arXiv preprint arXiv:2106.13304*.
- Patti, T. L., Kossaifi, J., Yelin, S. F., and Anandkumar, A. (2021b). Tensorly-quantum: Quantum machine learning with tensor methods. *arXiv preprint arXiv:2112.10239*.
- Peng, Y., Choi, B., and Xu, J. (2021). Graph learning for combinatorial optimization: A survey of state-of-the-art. *Data Science and Engineering*, 6(2):119–141.
- Pfeifer, R. N., Haegeman, J., and Verstraete, F. (2014). Faster identification of optimal contraction sequences for tensor networks. *Physical Review E*, 90(3):033315.
- Raffin, A., Hill, A., Ernestus, M., Gleave, A., Kanervisto, A., and Dormann, N. (2019). Stable baselines3. *GitHub repository*.
- Ran, S.-J., Tirrito, E., Peng, C., Chen, X., Tagliacozzo, L., Su, G., and Lewenstein, M. (2020). *Tensor Network Contractions: Methods and Applications to Quantum Many-Body Systems*. Springer Nature.
- Sato, R., Yamada, M., and Kashima, H. (2019). Approximation ratios of graph neural networks for combinatorial problems. *arXiv preprint arXiv:1905.10261*.
- Schindler, F. and Jermyn, A. S. (2020). Algorithms for tensor network contraction ordering. *Machine Learning: Science and Technology*, 1(3):035001.
- Schlag, S., Henne, V., Heuer, T., Meyerhenke, H., Sanders, P., and Schulz, C. (2016). K-way hypergraph partitioning via n-level recursive bisection. In *2016 Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 53–67. SIAM.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- 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):354–359.
- Singh, K. and Xie, M. (2008). Bootstrap: a statistical method. *Unpublished manuscript, Rutgers University, USA. Retrieved from http://www.stat.rutgers.edu/home/mxie/RCPapers/bootstrap.pdf*, pages 1–14.
- Stoudenmire, E. M. and Schwab, D. J. (2016). Supervised learning with quantum-inspired tensor networks. *arXiv preprint arXiv:1605.05775*.
- Xu, J., Liang, L., Deng, L., Wen, C., Xie, Y., and Li, G. (2019). Towards a polynomial algorithm for optimal contraction sequence of tensor networks from trees. *Physical Review E*, 100(4):043309.
- Yang, Y. and Whinston, A. (2020). A survey on reinforcement learning for combinatorial optimization. *arXiv preprint arXiv:2008.12248*.
- Yehudai, G., Fetaya, E., Meir, E., Chechik, G., and Maron, H. (2021). From local structures to size generalization in graph neural networks. In *International Conference on Machine Learning*, pages 11975–11986. PMLR.
- Zhao, L. and Akoglu, L. (2019). Pairnorm: Tackling over-smoothing in gnn. *arXiv preprint arXiv:1909.12223*.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., and Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81.
- Zhuang, V. and Sui, Y. (2021). No-regret reinforcement learning with heavy-tailed rewards. In *International Conference on Artificial Intelligence and Statistics*, pages 3385–3393. PMLR.

A. Additional Experimental details

A.1. Hyperparameters

Table 5 lists the hyperparameter values that were used for our trainable modules. Both the actor and the critic shared the same hyperparameters. Additionally, we scaled the rewards (FLOPS) by 10^{14} for the synthetic networks and the Sycamore circuit $m = 10, 12, 14$, while for the max-cut and sycamore and M20 circuits we used 10^{19} as our scaling factor.

Category	Parameter	synthetic networks	real-world networks
Random TNs	mean connectivity	3	-
Random TNs	tensor extent probability	$U[2, 6]$	-
GNN parameters	GNN layers	3	4
GNN parameter	# intermediate layer features	128	256
GNN parameter	Aggregation method	mean	mean
GNN parameter	Use Bias	False	False
RL	roll-out steps	4192	2096
RL	Training samples	1M	3M
RL	batch size	512	256
RL	Value weight	0.1	0.001
RL	Score Normalization, ϵ	10^{-2}	10^{-2}
RL	# paths during inference per inference step	50, 1000	160
RL	Inference step frequency	-	40
RL	PPO training step per roll-out collection	8	8
RL	Learning rate	$3 \cdot 10^{-2}$	$3 \cdot 10^{-2}$
RL	Entropy weight	0	0
RL	Optimistic buffer size	4096	4096

Table 5. Hyperparameters table

A.2. Wall-time Comparison

Table 6 compares the total cost of finding the optimal contraction path and the actual time of executing the contraction for different algorithms and networks. Recall that for QC applications, each network has to be contracted multiple times to get a high-fidelity estimate of the circuit output (Peng et al., 2021). Here, we assume that we contract each network $\ell = 10^6$ times (Pan et al., 2021b). The results indicate that RL-TNCO is 90%-40% faster than the best baseline considered.

network	RL-TNCO				Best baseline				RL-TNCO to Best baseline
	path optim' (sec)	contract. cost (flops)	single contract. (sec)	total (sec)	path optim' (sec)	contract. cost (flops)	single contract. (sec)	total (sec)	
max width= Random 100	1.6E+02	5.5E+12	2.7E-01	2.7E+05	1.0E+04	5.4E+13	2.7E+00	2.7E+06	0.10
Sycamore m=14	4.3E+04	2.6E+14	1.3E+01	1.3E+07	1.0E+04	5.0E+14	2.5E+01	2.5E+07	0.52
Sycamore m=20	4.3E+04	3.5E+18	1.7E+05	1.7E+11	1.0E+04	5.8E+18	2.9E+05	2.9E+11	0.60
MaxCut	8.6E+04	9.2E+12	4.6E-01	5.5E+05	2.1E+04	3.1E+13	1.5E+00	1.5E+06	0.35

Table 6. Wall time Comparison of path ordering and actual contraction time. Time is measured in seconds, cost in flops. Conversion is based on DGXA100 specifications.

A.3. Runtime vs performance trade-off: The Sycamore M20 circuit

In this subsection, we present a detailed performance comparison of RL-TNCO on the largest Sycamore circuit, Sycamore M20. A recent experiment on this circuit is claimed to prove quantum supremacy (Arute et al., 2019).

		Number of seeds								
		2	3	4	5	6	7	8	9	10
Limit (hrs)	0.5	7.54	7.53	7.475	7.42	7.42	7.42	7.14	7.14	7.14
	1	7.53	7.53	7.48	7.48	7.48	7.48	6.01	6.01	6.01
	2	6.7	6.45	6.43	6.43	6.43	6.43	5.83	5.83	5.83
	3	6.7	6.45	6.43	6.43	6.43	6.43	5.83	5.83	5.83

Table 7. The median contraction cost (in 10^{18} flops) for Sycamore M20 circuit using Contegra-Kayhpar. The estimation is based on the bootstrap method. For the 3 hours time limit runs, we present the best result out of the 2 hour run and 3 hour run

		Number of seeds								
		2	3	4	5	6	7	8	9	10
Limit (hrs)	0.5	5864	93.1	13.13	13.13	13.13	13.13	9.79	9.79	9.79
	1	5864	93.1	9.79	9.79	9.79	9.79	6.81	6.81	6.81
	2	127.5	93.1	7.38	7.38	7.38	7.38	5.25	5.25	5.25
	3	37.63	6.68	5.25	5.25	5.25	5.25	3.49	3.49	3.49

Table 8. The median contraction cost (in 10^{18} flops) for Sycamore M20 circuit using RL-TNCO. The estimation is based on the bootstrap method. Configuration which achieves superior performance over the corresponding setup using Cotengra-Kayhpar are highlighted in green, otherwise the cell is highlighted in red. The optimal result can be achieved with 8 seeds trained for three hours.

Optimizing time-constrained performance. When learning on a single network, one can improve performance by either using multiple seeds, resulting in different trained models or increasing the training time⁷. Likewise, other solvers may benefit from using multiple seeds, or extending the evaluation time. Our goal in this section is to investigate the performance gain when the the time budget is extended. This would allow us to understand what the optimal configuration is for a given time budget and in which configuration extending the solver time limit will result in decreasing marginal utility.

Both RL-TNCO and Cotengra-Kahypar are not deterministic, and we find that results tend to vary considerably across different initializations (different seeds). It therefore makes sense to run multiple instances with different seeds and select the best result. We are interested in estimating the result of each baseline with time limit τ and k seeds. Since this is an optimization score, the baseline score is the best score over the k seeds. The time limits ranged from 30 minutes to 3 hours, and we used up to 10 different seeds for each baseline.

Analysis. We used the bootstrap analysis method (Singh and Xie, 2008) to evaluate the median result. We run this baseline with 10 seeds for each time limit τ and obtained 10 results. We sampled with replacement $n \cdot k$ results, where $k = 2..10$, and obtained n tuples of size k . For each tuple $x_i(\tau, k)$ we took the best result, yielding n values $z_i(\tau, k)$. We report the median $m(\tau, k)$ over these values $\{z_i(\tau, k) | i = 1..n\}$ for Contegra Kahypar in Table 7 and for RL-TNCO in Table 8. In this graph, the cost of calculating the path order is negligible compared to the contraction time, and the gain from reducing contraction cost is by far higher than longer running times (Table ??).

A.4. Ablation study

To understand the contribution of components of our model, we performed an ablation study for the four components discussed in Section 6. Table 9 presents the results of the ablation study. The “Full” model is often the best configuration. Each component of RL-TNCO seems to improve the final result. Integrating existing solver information seems to have the most pronounced effect. Path pruning trades critic accuracy for train time. Fig 7 shows that a good result is achieved faster with path pruning.

A.5. Size generalization

The question whether GNNs trained on small graph can is an active research field (Yehudai et al., 2021)). In general, Performance is expected to deteriorate as the test distribution drifts from the train distribution. In our case, We ran the experiment proposed by the anonymous reviewers, and found that the effect was not pronounced. The performance on

⁷In setup 2: Learning and generalizing to new networks, it is also possible to increase the number of sampled trajectories.

	Full	w/o robust features	w/o solver info	w/o optimistic buffer	w/o path pruning
random TN 150 nodes	$\times 44.3$	$\times 12.6$	$\times 4.5$	$\times 6.1$	N/A
max-cut, last 100	$\times 6.2$	$\times 3.6$	$\times 0.7$	$\times 6.2$	$\times 6.2$
m12	$\times 7.2$	$\times 3.3$	$\times 10^{-3}$	$\times 6.0$	$\times 6.9$

Table 9. Improvement factor (best baseline FLOPS/RL-TNCO FLOPS) over best baseline for ablated models on various TNs. Path pruning during training is applicable in single network mode only. The max-cut results are for the suffix TN, excluding the shared prefix TN.

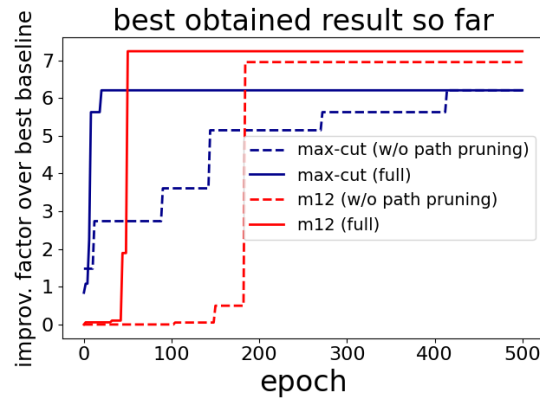


Figure 7. In a single-network setup we keep the best result obtained during training. This plot show the best improvement factor (best baseline FLOPS/FLOPS) as a function of the training epoch. Path pruning allows RL-TNCO to achieve a good result faster.

Optimizing Tensor Network Contraction Using Reinforcement Learning

the test set of 25 nodes TNs using models trained on TNs with 50, 75 and 100 nodes was 99.7%, 94.1%, 91.8% of the performance of models trained on TNs of size 25.