

# MARCO: A Memory-Augmented Reinforcement Framework for Combinatorial Optimization

Andoni I. Garmendia<sup>1</sup>, Quentin Cappart<sup>2</sup>, Josu Ceberio<sup>1</sup> and Alexander Mendiburu<sup>1</sup>

<sup>1</sup>University of the Basque Country (UPV/EHU), Donostia-San Sebastian, Spain

<sup>2</sup>Polytechnique Montréal, Montreal, Canada

{andoni.irazusta, josu.ceberio, alexander.mendiburu}@ehu.eus, quentin.cappart@polymtl.ca

## Abstract

*Neural Combinatorial Optimization* (NCO) is an emerging domain where deep learning techniques are employed to address combinatorial optimization problems as a standalone solver. Despite their potential, existing NCO methods often suffer from inefficient search space exploration, frequently leading to local optima entrapment or redundant exploration of previously visited states. This paper introduces a versatile framework, referred to as *Memory-Augmented Reinforcement for Combinatorial Optimization* (MARCO), that can be used to enhance both constructive and improvement methods in NCO through an innovative memory module. MARCO stores data collected throughout the optimization trajectory and retrieves contextually relevant information at each state. This way, the search is guided by two competing criteria: making the best decision in terms of the quality of the solution and avoiding revisiting already explored solutions. This approach promotes a more efficient use of the available optimization budget. Moreover, thanks to the parallel nature of NCO models, several search threads can run simultaneously, all sharing the same memory module, enabling an efficient collaborative exploration. Empirical evaluations, carried out on the maximum cut, maximum independent set and travelling salesman problems, reveal that the memory module effectively increases the exploration, enabling the model to discover diverse, higher-quality solutions. MARCO achieves good performance in a low computational cost, establishing a promising new direction in the field of NCO.

## 1 Introduction

The objective in Combinatorial Optimization (CO) problems is to find the optimal solution from a finite or countable infinite set of discrete choices. These problems are prevalent in many real-world applications, such as chip design [Mirhoseini *et al.*, 2021], genome reconstruction [Vrček *et al.*, 2022] and program execution [Gagrani *et al.*, 2022].

In recent years, the field of *Neural Combinatorial Optimization* (NCO) has emerged as an alternative tool for solving such problems [Bengio *et al.*, 2021; Mazyavkina *et al.*, 2021; Bello *et al.*, 2016]. NCO uses deep neural networks to address CO problems in an end-to-end manner, learning from data and generalizing to new, unseen instances. Researchers in this field have followed the steps of heuristic optimization, proposing the neural counterparts of *constructive methods* [Bello *et al.*, 2016; Kool *et al.*, 2018; Kwon *et al.*, 2020] and *improvement methods* [Lu *et al.*, 2019; Chen and Tian, 2019; Wu *et al.*, 2021].

*Neural constructive methods* quickly generate an approximate solution in a one-shot manner by means of a learnt neural model. While being simple and direct, constructive methods suffer from their irreversible nature, barring the possibility of revisiting earlier decisions. This limitation becomes particularly pronounced in large problems where suboptimal initial decisions in the construction of the solution can significantly impact the final outcome. To improve the performance of these methods, recent efforts have employed techniques such as *sampling*, where instead of following the output of the model deterministically, a random sample is taken from a probability distribution given by the output, with the intention of obtaining better solutions and break with the deterministic behaviour, obtaining a richer set of solutions; or *beam search* [Choo *et al.*, 2022], which maintains a collection of the highest-quality solutions as it explores the search space based on the output of the neural network, i.e., the probability of adding an item to the partial solution that is being constructed. Similarly, *active search* [Bello *et al.*, 2016; Hottung *et al.*, 2021] is used to update the model’s weights (or a particular set of weights) during test time, in order to overfit the model to the test instance to be solved.

Alternatively, *neural improvement methods* are closely linked to perturbation methods, such as local search. They start from a complete solution, and operate by iteratively suggesting a modification that improves the current solution at the present state. Unlike constructive methods, improvement methods inherently possess the ability to explore the search space of complete solutions. However, they often get stuck in local optima or revisit the same states repeatedly, leading to cyclical patterns. Recent studies [Barrett *et al.*, 2020; Garmendia *et al.*, 2023] have employed a variety of strategies inherited from the combinatorial optimization literature

to tackle these drawbacks. The method by [Barrett *et al.*, 2020] keeps a record of previously performed actions, while the study in [Garmendia *et al.*, 2023] maintains a tabu memory of previously visited states, forbidding the actions that would lead to visit those states again.

Neural constructive methods, neural improvement methods, and most classical optimization proposals all face a significant challenge: *exploring efficiently the search space*. To address this, we introduce a new framework, referred to as *Memory-Augmented Reinforcement for Combinatorial Optimization*, or MARCO. This framework integrates a memory module into both neural constructive and neural improvement methods. The memory records the visited or created solutions during the optimization process, and retrieves relevant historical data directly into the NCO model, enabling it to make more informed decisions.

A key feature of MARCO is the ability to manage a shared memory when several search *threads* are run in parallel. By doing so, MARCO not only reduces the redundancy of storing similar data across multiple threads but also facilitates a collaborative exploration of the search space, where each thread benefits from a collective understanding of the instance.

The main contributions of the paper are as follows: (1) introducing MARCO as a pioneering effort in integrating memory modules within both neural improvement and constructive methods. (2) Designing a similarity-based search mechanism that retrieves past, relevant information to feed the memory and to better inform the model. (3) Presenting the parallelism capabilities of MARCO, which enables a more efficient and collaborative exploration process. (4) Illustrating the implementation of the framework to three graph-based problems: maximum cut, maximum independent set, and travelling salesman problem. Experiments are then carried out on these three problems with graphs up to 1200 nodes. The empirical results indicate that MARCO surpasses some of the recently proposed learning-based approaches, demonstrating the benefits of using information regarding visited solutions. The source code and supplementary material are available online<sup>1</sup>.

## 2 Related Work

Various strategies have been developed to enhance the exploration of the search space in NCO algorithms. Most of the methods sample from the model’s logits [Bello *et al.*, 2016; Kool *et al.*, 2018; Kwon *et al.*, 2020], which introduces stochasticity into the solution inference process. Beyond sampling, entropy regularization has been implemented during the training of NCO models [Kim *et al.*, 2021], to ensure the models are not overconfident in their output. Furthermore, [Grinsztajn *et al.*, 2024] proposed a multi-decoder system, where each decoder is trained on instances where it performs best, resulting in a set of specialized and complementary policies.

Despite these advancements, none of these methods exploit any kind of memory mechanism, which has the potential to

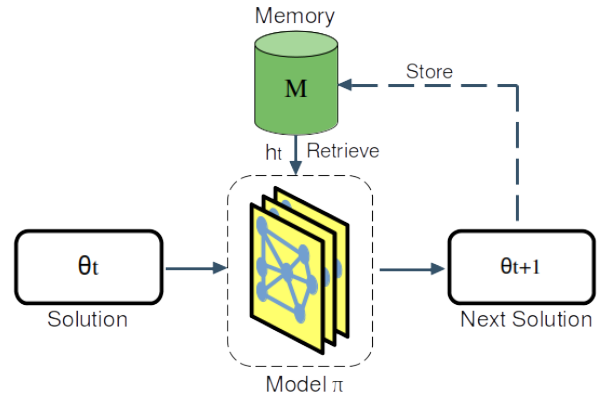


Figure 1: Optimization step in MARCO.

leverage previous experiences in the decision-making process and promote exploration.

In the work by [Garmendia *et al.*, 2023], a *tabu search* algorithm [Glover and Taillard, 1993], known for its memory-based approach to circumventing cyclical search patterns, is layered on top of a neural improvement method. The algorithm utilizes a tabu memory to track previously visited solutions. However, this memory serves merely as an external filter, preventing the selection of tabu actions without integrating historical data into the neural model’s decision-making process.

DeepACO [Ye *et al.*, 2023] uses a neural network to learn the underlying heuristic of an *ant colony optimization* algorithm [Blum, 2005; Dorigo *et al.*, 2006]. It maintains an external pheromone matrix, indicative of promising variable decisions. However, the integration of this pheromone data is indirect; it is combined in a post-hoc fashion with the output probabilities of the model rather than being an intrinsic part of the learning process.

Closer to our work, ECO-DQN [Barrett *et al.*, 2020] is a neural improvement method that records the last occurrence of each action. This operation-based memory approach, which simply tracks when actions were last taken, is computationally efficient, requiring only minimal storage. The drawback of this approach is that it only focuses on the actions, failing to consider the overall search context. The effectiveness of an action is often contingent on the broader state of the optimization process, a fact that operation-based memory fails to capture. Compared to this work, we save entire solutions in memory, incorporating a more holistic view of the search context to the system, at the cost of higher memory requirements.

## 3 MARCO: A Memory-Based Framework

This section introduces MARCO, the main contribution of this paper. Although the framework can be used for arbitrary CO problems, we first focus on *graph-based problems*, as they are ubiquitous in combinatorial optimization. In fact, from the 21 NP-complete problems identified by Karp (1972), ten are decision versions of graph optimization problems, while most of the other ones can also be modeled over graphs.

<sup>1</sup><https://github.com/TheLeprechaun25/MARCO>.

Let  $G = (V, E)$  be a simple graph composed of a set of nodes  $V$  and a set of edges  $E$ . Finding a solution  $\theta$  for graph problems often involves finding subsets of nodes or edges that satisfy specific criteria, such as minimizing or maximizing a certain objective function.

Briefly, the idea of MARCO is to leverage both (1) a learnt policy  $\pi$  defining how the current solution should be modified for exploring the search space, and (2) a memory module  $\mathcal{M}$ , providing information to build the policy. The policy is typically parameterized with a neural network, and especially with a *graph neural network* [Kipf and Welling, 2016] when operating on graph problems. Such an architecture has been considered as highly relevant for combinatorial optimization [Cappart *et al.*, 2023]. Besides, the policy is iteratively called to modify the solution until a convergence threshold has been reached. A typical optimization step with MARCO is illustrated in Figure 1.

This mechanism can be integrated into both constructive and improvement methods. The main difference relates to *how a solution is defined* and *how information is retrieved from the memory*. Let  $\theta_t$  refer to a complete solution obtained after  $t$  iterations, and  $\hat{\theta}_t$  refer to a partial solution, i.e., a solution where only  $t$  variables has been assigned, with at least one variable not assigned. In *constructive methods*, MARCO is capable of using a deterministic policy repeatedly, i.e., opting for the greedy action to generate multiple different constructions. Each construction starts from an empty solution and each optimization step consists in extending the current partial solution, i.e., assigning an unassigned variable in the optimization problem. The policy takes as input both static information (i.e., the graph instance  $G$ ) and dynamic information related to the current partial solution ( $\hat{\theta}_t$ ). The memory then stores a solution once it is completed, i.e., once all the variables have been assigned. On the other hand, *improvement methods* feature an iterative refinement of a complete solution. Each step modifies a current (complete) solution  $\theta_t$ , transitioning it to a subsequent solution  $\theta_{t+1}$ . In this scenario, the dynamic information is the complete solution  $\theta_t$ . Each explored solution is recorded into the memory.

For both methods, the training is conducted through reinforcement learning. Each time a completed solution is reached, a reward  $r_t$  is obtained, denoting how good the executed optimization trajectory has been. The reward is designed to balance two factors: (1) the quality of the solution found and (2) the dissimilarity of the new solution compared to previous solutions stored in memory.

### 3.1 Memory Module

As shown in Algorithm 1, the inference in MARCO starts with the selection of an initial solution (refer to line 2). In each optimization step, the memory module  $\mathcal{M}$  is responsible for storing the visited solutions (line 6), and retrieving aggregated historical data  $h_t$  (line 7). The historical data ( $h_t$ ) is aggregated with the current (partial) solution and the graph features ( $G$ ) to form the current optimization state  $s_t$  (line 8). Subsequently,  $s_t$  is input into the model (line 9), which then proposes a set of actions that generate new solutions (line 10).

The specific process of retrieval is shown in Algorithm 2. To retrieve relevant solutions, MARCO employs a *similarity-*

*based search*. This involves comparing the current (partial) solution ( $\hat{\theta}_t$  or  $\theta_t$ ) with each stored solution ( $\theta_{t'}$  where  $t' < t$ ) using a similarity metric (e.g., the inner product in line 4). Intuitively, the idea is to feed the policy with the most similar solutions to the current one for executing the next exploration step. We carry out the retrieval using a *k-nearest neighbors search* (line 5). Rather than simply averaging the  $k$  most similar solutions, MARCO uses a weighted average approach, where the weight given to each past solution is directly proportional to its similarity to the current solution. This score is normalized, ranging from 0 (completely different) to 1 (identical), to represent the level of similarity (see line 6).

---

#### Algorithm 1 Inference with MARCO

---

```

1: procedure MARCO(graph  $G$ , policy  $\pi$ ,  $k$ ,  $max\_steps$ )
2:    $\theta_0 \leftarrow$  INITIALIZESOLUTIONS
3:    $\mathcal{M} \leftarrow$  INITIALIZEMEMORY( $k$ )
4:   for  $t = 0$  to  $max\_steps - 1$  do
5:     if  $\theta_t$  is Completed then
6:        $\mathcal{M} \leftarrow$  STOREINMEM( $\theta_t$ )
7:        $h_t \leftarrow$  RETRIEVEFROMMEM( $\mathcal{M}$ ,  $k$ ,  $\theta_t$ )
8:        $s_t \leftarrow$  AGGREGATE( $G$ ,  $\theta_t$ ,  $h_t$ )  $\triangleright$  Get current state
9:        $a_t \leftarrow$  POLICY( $\pi$ ,  $s_t$ )
10:       $\theta_{t+1} \leftarrow$  STEP( $\theta_t$ ,  $a_t$ )  $\triangleright$  Get next solution

```

---



---

#### Algorithm 2 Action retrieval from memory

---

```

1: procedure RETRIEVEFROMMEM( $\mathcal{M}$ ,  $k$ ,  $\theta_t$ )
2:    $v \leftarrow$  LENGTH( $\mathcal{M}$ )
3:    $k \leftarrow$  min( $k$ ,  $v$ )
4:    $simScore \leftarrow$  INNERPRODUCT( $\theta_t$ ,  $\theta_{t'} \mid t' < t$ .)
5:    $h_t \leftarrow$  KNN( $k$ ,  $simScore$ ,  $\mathcal{M}$ )
6:    $\hat{h}_t \leftarrow$   $h_t \times$  NORM( $simScore$ )
7:   return  $\hat{h}_t$   $\triangleright$  Return relevant historical data

```

---

### Collaborative Memory

An additional feature enabled by MARCO is the implementation of parallel optimization threads during its inference phase. In this setup, multiple concurrent threads are run for each problem instance, collaboratively exploring the solution space. A key aspect of this functionality is the use of shared memory across all threads. This collective memory stores all the explored solutions by any thread, making it accessible to the entire group.

## 4 Application of MARCO

In this study, we demonstrate the adaptability of MARCO to various problem types, encompassing both constructive and improvement methods. We specifically apply MARCO in two scenarios: (1) a neural improvement method for problems with binary variables, such as the Maximum Cut (MC) and the Maximum Independent Set problem (MIS); and (2) a neural constructive method for permutation problems, such as the Travelling Salesman Problem (TSP).




Figure 2: MARCO for Neural Improvement methods for the Maximum Cut problem. Initially, multiple solutions are randomly generated for a given problem instance. Each solution is iteratively improved, forming a thread. Throughout this process, the visited solutions and corresponding actions are stored into a shared memory. This collective memory then updates the graph features fed to the model.

#### 4.1 Improvement Methods for Binary Problems

In binary optimization problems, a solution is formalized as a binary vector, denoted as  $\theta \in \{0, 1\}^{|V|}$  for a problem with  $|V|$  variables. Each variable  $x_i$  represents a binary decision for the  $i^{\text{th}}$  variable. Neural improvement methods are designed to optimize a problem by iteratively refining an initial complete solution  $\theta_0$ , which can be generated either randomly or through heuristic methods. In the case of binary problems, the central operation is a node-wise operator that flips the value of a node in  $\theta$ . The memory records visited solutions and their associated actions (e.g., a *bit-flip* action, consisting in flipping the value of a variable). When a new solution is generated, the model consults the memory to retrieve the actions performed in similar previous scenarios. The importance of the stored actions is given by the similarity between the current solution  $\theta_t$  and previously stored solutions  $\theta_{t'}$  with  $t' < t$ . We compute the similarity using the inner product:

$$\text{Similarity}(\theta_t, \theta_{t'}) = \langle \theta_t, \theta_{t'} \rangle = \sum_{i \in V} (\theta_t)_i \cdot (\theta_{t'})_i \quad (1)$$

In this case, the aggregated memory data ( $h_t$ ) is a vector of size  $|V|$ , defined as the weighted average of the actions that were executed in the  $k$  most similar solutions (if any). See Figure 2 for a visual description of the inference in neural improvement methods with MARCO.

The *reward*  $r_t$  obtained by a neural improvement model at each step  $t$  is defined as the non-negative difference between the current objective value of the solution,  $f(\theta_t)$ , and the best objective value found thus far ( $f(\theta^*)$ ), i.e.,  $r_t = \max\{f(\theta_t) - f(\theta^*), 0\}$ . This reward structure, prevalent in neural improvement methods [Ma *et al.*, 2021; Wu *et al.*, 2021], motivates the model to continually seek better solutions. To prevent the model from cycling through the same states and encourage novel solution exploration, we incorporate a binary penalty term  $p_r$ , activated when revisiting previously encountered solutions. The adjusted reward for each




Figure 3: MARCO for Neural Constructive methods. Travelling Salesman example. Each solution in a batch begins with a distinct initial node. Subsequently, every thread proceeds to iteratively construct a solution, considering data gathered from the memory module. Upon the completion of this construction process, the obtained solution is stored within the memory, serving as a reference for subsequent solution constructions.

step is thus  $\hat{r}_t = r_t - w_p \times p_r$ , where  $w_p$  is a weight factor for the penalty.

#### 4.2 Constructive Methods for Permutations

The objective in permutation problems like the TSP is to find a permutation of nodes in a graph that maximizes or minimizes a specific objective function. Neural constructive methods build the permutation incrementally, starting from an empty solution and adding elements sequentially until a complete permutation is formed.

In the context of permutation problems, the solution  $\theta$  can also be conceptualized as a binary vector  $\theta^b \in \{0, 1\}^{|E|}$ . Each element in this vector corresponds to an edge  $e_{ij}$  in the graph, indicating whether the edge is part of the solution, that is, whether node  $i$  and node  $j$  are adjacent in the permutation.

At each step of the permutation building process, the model operates on a partial solution, defined as a sequence  $\hat{\theta}_t^b = (\theta_t^b[1], \theta_t^b[2], \dots, \theta_t^b[k])$ , where  $k < |V|$  is the current number of nodes in the sequence. As the model progresses through constructing the permutation, the memory data is used to consider which edges have been selected in previously constructed solutions that are similar to  $\hat{\theta}_t^b$ . The similarity score is performed by an inner product between the binary representations of the partial solution  $\hat{\theta}_t^b$  and the complete solutions saved in memory  $\theta_{t'}^b$  with  $t' < t$ :

$$\text{Similarity}(\hat{\theta}_t^b, \theta_{t'}^b) = \langle \hat{\theta}_t^b, \theta_{t'}^b \rangle = \sum_{i \in V} (\hat{\theta}_t^b)_i \cdot (\theta_{t'}^b)_i \quad (2)$$

Figure 3 showcases the inference in MARCO for neural constructive methods. Training involves computing a *reward* once the solution is completed. The reward  $r_t = f(\theta_t)$ , given by the objective value of the solution, is adjusted by subtracting a baseline value to stabilize training. A common approach is to use the average reward across different initializations, as done in POMO for the TSP [Kwon *et al.*, 2020].

Our initial experiments with constructive models showed that exact solution repetitions are uncommon for large instances. Therefore, instead of the binary penalty system used in improvement methods, we apply a scaled penalization based on similarity levels with stored solutions. The final reward is calculated as  $\hat{r}_t = r_t - w_p \times \text{AVGSIM}(\theta_t, \theta_{t'})$ , where  $\text{AVGSIM}(\theta_t, \theta_{t'})$  is the average of all the inner products between the constructed solution and the  $k$  most similar stored solutions.

## 5 Model Architecture

Graph neural networks are particularly well-suited to parameterize policy  $\pi$ . We specifically use a Graph Transformer (GT) [Dwivedi and Bresson, 2020] coupled with a Feed-Forward Neural Network. GTs are a generalization of transformers [Vaswani *et al.*, 2017] for graphs. The fundamental operation in GTs involves applying a shared self-attention mechanism in a fully connected graph, allowing each node to gather information from every other node. The gathered information is then weighted by computed attention weights, which indicate the importance of each neighbor’s features to the corresponding node.

Our model aims to be adaptable to various combinatorial problems, requiring it to assimilate both the graph’s structural information and the attributes of its nodes and edges. To achieve this, we modify the GT to incorporate structural information encoded as edge features within the Attention (Attn) mechanism. This adaptation is reflected in the following equation.

$$\text{Attn}(Q, K, V) = \left( \text{softmax} \left( \frac{QK^T}{\sqrt{d_k}} + \mathbf{E} \right) \cdot \mathbf{E} \right) V \quad (3)$$

In this equation,  $Q$ ,  $K$ , and  $V$  stand for Query, Key, and Value, respectively, which are fundamental components of the attention mechanism [Vaswani *et al.*, 2017] and  $d_k$  is a scaling factor.  $\mathbf{E} = \mathbf{W}_e \cdot \mathbf{e}_{ij}$  is a linear transformation of the edge weights, where  $\mathbf{W}_e \in \mathbb{R}^{1 \times n_{\text{heads}}}$  is a learnable weight matrix, and  $\mathbf{e}_{ij}$  represents the edge features between nodes  $i$  and  $j$ .  $\mathbf{E}$  integrates edge information by being added to the attention scores and used in a dot product.

The final step involves processing the output of the GT through an element-wise feed-forward neural network to generate action probabilities. The output of the GT could be both node- or edge-embeddings. The performed action depends on the method in use. In our studied cases, we will use node embeddings to generate node-based actions: node-flips for improvement methods in binary problems and node addition to the partial solution for the constructive method in permutation problems. However, MARCO is also applicable to edge-based actions, such as pairwise operators (swap, 2-opt) for permutation-based improvement methods.

We utilize the policy gradient REINFORCE algorithm [Williams, 1992] to find the optimal parameter set,  $\pi^*$ , which maximizes the expected cumulative reward in the optimization process.

## 6 Experiments

### 6.1 Problems

We validate the effectiveness of MARCO across a diverse set of CO problems both binary and permutation-based: the Maximum Cut (MC), Maximum Independent Set (MIS) and the Travelling Salesman Problem (TSP).

**Maximum Cut (MC).** The objective in MC [Dunning *et al.*, 2018] is to partition the set of nodes  $V$  in a graph  $G$  into two disjoint subsets  $V_1$  and  $V_2$  such that the number of edges between these subsets (the cut) is maximized. The objective function can be expressed as:  $\max \sum_{(u,v) \in E} \delta_{[\theta_u \neq \theta_v]}$  where  $\theta_u$  and  $\theta_v$  are binary variables indicating the subset to which nodes  $u$  and  $v$  belong, and  $\delta$  is a function which equals to 1 if  $\theta_u$  and  $\theta_v$  are different and 0 otherwise.

**Maximum Independent Set (MIS).** For the MIS problem [Lawler *et al.*, 1980], the goal is to find a binary vector  $\theta$  that represents a subset of nodes  $S \subseteq V$  in a graph  $G$  such that no two nodes in  $S$  are adjacent, and the size of  $S$  is maximized. The objective function can be formulated as:  $\max |S|$  such that  $(u, v) \notin E$  for all  $u, v \in S$

**Travelling Salesman Problem (TSP).** In TSP [Lawler *et al.*, 1986; Wang and Tang, 2021], given a set of nodes  $V$  and distances  $d_{u,v}$  between each pair of nodes  $u, v \in V$ , the task is to find a permutation  $\theta$  of nodes in  $V$  that minimizes the total travel distance. This is expressed as:  $\min \sum_{i=1}^{|V|} d(\theta_i, \theta_{i+1})$  with  $\theta_{|V|+1} = \theta_1$

### 6.2 Experimental Setup

**Training** For each problem, we train a unique model, using instances that vary in size. This helps the model to learn strategies that can be transferable between differently sized instances. For the MC and MIS, we used randomly generated Erdos-Renyi (ER) [Erdős *et al.*, 1960] graphs with 15% of edge probability, and sizes ranging from 50 to 200 nodes. For the TSP, fully connected graphs ranging from 50 to 100 nodes were generated, in which cities were sampled uniformly in a unit square. The total training time depends on the problem. The models for both MC and MIS required less than 40 minutes, while the one for the TSP required a significantly longer training (4 days) to reach convergence. See the supplementary material for a detailed description of the training configuration used.

**Inference** To evaluate the performance of MARCO, we have established certain inference parameters. For MC and MIS, we set the neural improvement methods to execute with 50 parallel threads (processing 50 solutions simultaneously), stopping upon  $2|V|$  improvement steps. For the TSP, we use 100 parallel initializations (as done in POMO [Kwon *et al.*, 2020]) and 20 iterations (solution constructions) for each instance. We have used  $k = 20$  for the similarity search. A more detailed description of the inference configuration used is reported in the supplementary material. MARCO has been implemented using *PyTorch 2.0*. A *Nvidia A100* GPU has been used to train the models and perform inference. Exact methods and heuristics serving as baselines were executed in a cluster with *Intel Xeon X5650* CPUs.

**Evaluation Data** Following the experimental setup of recent works [Ahn *et al.*, 2020; Böther *et al.*, 2021; Zhang *et al.*, 2023], we will evaluate the MC and MIS problems in ER graphs of sizes between 700-800, and harder graphs from the RB benchmark [Xu and Li, 2000] of sizes between 200-300 and 800-1200. For TSP, we follow the setting from [Kool *et al.*, 2018] and use randomly generated instances, with uniformly sampled cities in the unit square. We use graphs of sizes 100, 200 and 500.

**Ablations** We evaluate MARCO through several ablations that help us understand the impact of its different components. We begin by evaluating standalone models proposed in this work: the Neural Improvement Method (NIM) and Neural Constructive Methods (NCM), both of which operate without any integrated memory module. Next, for improvement methods, we add a NIM equipped with an operation-based memory (Op-NIM), tracking the number of steps since each action was executed lastly (imitating ECO-DQN [Barrett *et al.*, 2020]). Finally, we assess MARCO-ind, a variant of MARCO that operates without shared memory, executing multiple threads simultaneously but *independently*, with each thread maintaining its own separate memory.

**Baselines** To assess MARCO’s performance, we conduct a comprehensive comparison against a broad spectrum of combinatorial optimization methods tailored to each specific problem addressed. Our comparative analysis includes exact algorithms, heuristics, and learning-based approaches

For the MC, our comparison includes the GUROBI solver [Gurobi Optimization, LLC, 2023], the local search enhanced heuristic BURER [Burer *et al.*, 2002], and ECO-DQN [Barrett *et al.*, 2020], which is a neural improvement method incorporating an operation-based memory.

For MIS, we also include GUROBI [Gurobi Optimization, LLC, 2023], together with KAMIS [Lamm *et al.*, 2016], a specialized algorithm for MIS; and a constructive heuristic (Greedy), that selects the node with minimum degree in each step. Furthermore, we examine also a range of recently proposed learning-based methods: DGL [Böther *et al.*, 2021], LwD [Ahn *et al.*, 2020] and FlowNet [Zhang *et al.*, 2023].

For the TSP, we report results of the well known conventional solver Concorde [Applegate *et al.*, 2006], the heuristic LKH-3 [Papadimitriou, 1992], the Nearest Neighbor (NN) heuristic; and the learning-based methods used are the neural constructive POMO [Kwon *et al.*, 2020] enhanced with sampling and data augmentation, LEHD [Luo *et al.*, 2024] which reports the best results among neural methods and two of the state-of-the-art neural improvement methods: DACT [Ma *et al.*, 2021] and NeuOPT [Ma *et al.*, 2023].

### 6.3 Results

We present the results for each studied problem in a table divided by three row-segments, the first one consisting of non-learning methods (exact and heuristic), the second with recent learning methods from the literature, and the third with the methods (MARCO and ablations) proposed in this paper. We report both the average objective value in the evaluation instance set and the time needed for performing inference with a unique instance (batch size of 1). We use *ms*, *s* and *m* to

denote milliseconds, seconds and minutes, respectively. For learning methods, we report the results from the best performing configuration reported in the original paper. For exact solvers, we report the best found solution when the optimal solution is not achieved in a limit of 1 and 10 minutes per instance.

**MC** In Table 1 we report the results for the MC. MARCO significantly outperforms GUROBI and ECO-DQN, especially in larger problem instances (ER700-800, RB800-1200). In addition, MARCO proves to be competitive against the state-of-the-art heuristic, BURER, in the studied graph instances. The ablation results show that using the proposed memory scheme is superior to (1) not using any memory module, and (2) using an operation-based memory. Moreover, using a shared memory slightly improves the performance (with respect to MARCO-ind), while the computational cost is reduced. Compared to the ECO-DQN in computational cost, MARCO reduces the time needed to perform  $2|V|$  improvement steps.

Method	ER700-800		RB200-300		RB800-1200	
	Obj. $\uparrow$	Time	Obj. $\uparrow$	Time	Obj. $\uparrow$	Time
GUROBI	23420.17	1m	2024.55	1m	20290.08	1m
GUROBI <sub>long</sub>	24048.93	10m	2286.48	10m	23729.44	10m
BURER	<b>24235.93</b>	1.0m	<b>2519.47</b>	1.0m	<b>29791.52</b>	1.0m
ECO-DQN	24114.06	2.1m	2518.76	29s	29638.78	3.0m
NIM	24037.66	45s	2517.01	1.5s	29752.92	2.0m
Op-NIM	24081.18	47s	2518.34	1.6s	29751.87	2.1m
MARCO-ind	24203.11	52s	2519.46	2.3s	29778.84	2.7m
MARCO	<b>24205.97</b>	49s	<b>2519.47</b>	2.2s	<b>29780.71</b>	2.5m

Table 1: MC performance table. The best results overall and the best results among learning-based methods are highlighted in bold.

**MIS** Table 2 summarizes the results for MIS. Here, MARCO is also able to surpass the learning methods and its ablations, obtaining a comparable performance to the exact solver. Moreover, it reduces the gap to the specialized KAMIS algorithm. While incorporating a memory module in MARCO (NIM vs. MARCO) increases the time cost, it contributes to achieving superior solutions, while NIM gets stuck in suboptimal solutions (increasing the number of steps does not increase the performance).

**TSP** Results for the TSP are reported in Table 3. MARCO can obtain good inference performance in the studied instances, reaching the best found solutions for N100 and N200; and being second on N500, only surpassed by LEHD. It is important to note that our basic NCM implementation (without memory) obtains comparable results with the state-of-the-art learning method while being orders of magnitude faster. Also, MARCO improves over both NCM and the method without sharing memory (MARCO-ind).

### Generalization to Larger Sizes

Training NCO models with reinforcement learning is computationally intensive, leading to a common practice in the

Method	ER700-800		RB200-300		RB800-1200	
	Obj. $\uparrow$	Time	Obj. $\uparrow$	Time	Obj. $\uparrow$	Time
GUROBI	43.47	1m	19.98	1m	40.90	1m
GUROBI <sub>long</sub>	43.64	10m	20.03	10m	41.34	10m
KAMIS	<b>44.98</b>	1m	<b>20.10</b>	1m	<b>43.15</b>	1m
Greedy	38.85	50ms	18.41	4ms	37.78	54ms
DGL	38.71	11s	19.01	2s	32.32	3s
LwD	41.17	4s	17.36	1s	34.50	1s
FlowNet	41.14	2s	19.18	0.1s	37.48	0.5s
NIM	40.16	2s	19.26	0.5s	37.80	1s
Op-NIM	40.66	4s	19.70	1.2s	38.59	4s
MARCO-ind	43.72	19s	19.77	1.5s	39.94	7s
MARCO	<b>43.78</b>	17s	<b>19.87</b>	1.4s	<b>40.13</b>	6s

Table 2: MIS performance table. The best results overall and the best results among learning-based methods are highlighted in bold.

Method	N100		N200		N500	
	Obj. $\downarrow$	Time	Obj. $\downarrow$	Time	Obj. $\downarrow$	Time
Concorde	<b>7.76</b>	1m	<b>10.72</b>	1m	<b>16.59</b>	1m
LKH-3	<b>7.76</b>	1m	<b>10.72</b>	1m	<b>16.59</b>	1m
NN	9.69	1ms	13.45	2ms	20.80	5ms
POMO	7.81	0.1s	11.73	1s	21.88	2s
LEHD	<b>7.76</b>	2m	<b>10.72</b>	4m	<b>16.63</b>	10m
DACT	7.77	3m	14.23	5m	145.78	11m
NeuOPT	7.77	2m	10.73	4m	39.19	8m
NCM	<b>7.76</b>	0.1s	10.75	1s	16.90	2s
MARCO-ind	<b>7.76</b>	3s	10.73	11s	16.81	22s
MARCO	<b>7.76</b>	3s	<b>10.72</b>	11s	16.78	21s

Table 3: TSP performance table. The best results overall and the best results among learning-based methods are highlighted in bold.

literature where models are often trained on smaller-sized instances (up to 100). While this approach is understandable due to resource constraints, it is important to consider the ability of these models to generalize to larger instances. This aspect is crucial for their applicability in real-world scenarios, where problem sizes can vary significantly.

The data presented in Table 2 illustrate this point. Here, even a basic greedy constructive method manages to outperform more complex learning-based methods (DGL, LwD, and FlowNet) in the RB800-1200 instance. This observation underlines the importance of using simple heuristics as a sanity check to assess whether advanced models are effectively generalizing to unseen instances or larger sizes. Similarly, Table 3 reveals that a simple Nearest Neighbour heuristic is able to surpass POMO, DACT and NeuOPT in instances of 500 cities. Even though the underlying model of MARCO has been trained on smaller instances (up to 200 for MC and MIS, and up to 100 for TSP), it is able to maintain a good performance in larger graphs with a lower time cost compared to state-of-the-art heuristic solvers.

## 7 Limitations and Future Work

MARCO offers significant advancements in neural combinatorial optimization. However, it has room for improvement. A primary concern is the uncontrolled growth of its memory during the optimization process, as it continually stores all the encountered states, leading to increased computational and memory costs. To counter this, future work could focus on implementing mechanisms to prune the memory by removing redundant information.

Another limitation is the substantial resource requirement for storing entire edge-based solutions in memory (like in TSP). This approach, particularly for large instances, can result in high memory consumption and slower retrieval processes. A promising direction would be to represent solutions in a lower-dimensional space using fixed-size embeddings, effectively reducing the memory footprint while preserving (or even incorporating) necessary information.

In terms of data retrieval, MARCO currently employs a method based on a weighted average of similarity, which may not fully capture the relationships between solution pairs. A more advanced alternative to consider is the implementation of an attention-based search mechanism. This method would not only prioritize the significance of various stored solutions but could also incorporate the objective values or other distinct characteristics of these solutions to compute their relevance.

Additionally, while not a limitation, applying MARCO to new problems or integrating it with different NCO methods requires careful consideration in how memory information is aggregated with instance features. The nature of the data stored and retrieved can vary significantly depending on the specific problem being addressed.

## 8 Conclusion

In this paper, we have introduced the Memory-Augmented Reinforcement for Combinatorial Optimization (MARCO), a framework for Neural Combinatorial Optimization methods that employs a memory module to store and retrieve relevant historical data throughout the search process. The experiments conducted in the maximum cut, maximum independent set and travelling salesman problems validate MARCO’s ability to quickly find high-quality solutions, outperforming or matching the state-of-the-art learning methods. Furthermore, we have demonstrated that the use of a collaborative parallel-thread scheme contributes to the performance of the model while reducing the computation cost.

## Acknowledgments

Andoni Irazusta Garmendia acknowledges a predoctoral grant from the Basque Government (ref. PRE.2020.1\_0023). This work has been partially supported by the Research Groups 2022-2025 (IT1504-22), the Elkartek Program (KK-2021/00065, KK-2022/00106) from the Basque Government.

## References

[Ahn *et al.*, 2020] Sungsoo Ahn, Younggyo Seo, and Jinwoo Shin. Learning what to defer for maximum independent

- sets. In *International Conference on Machine Learning*, pages 134–144. PMLR, 2020.
- [Applegate *et al.*, 2006] David Applegate, Ribert Bixby, Vasek Chvatal, and William Cook. Concorde tsp solver, 2006. URL <http://www.tsp.gatech.edu/concorde>, 2006.
- [Barrett *et al.*, 2020] Thomas Barrett, William Clements, Jakob Foerster, and Alex Lvovsky. Exploratory combinatorial optimization with reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 3243–3250, 2020.
- [Bello *et al.*, 2016] Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*, 2016.
- [Bengio *et al.*, 2021] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.
- [Blum, 2005] Christian Blum. Ant colony optimization: Introduction and recent trends. *Physics of Life reviews*, 2(4):353–373, 2005.
- [Böther *et al.*, 2021] Maximilian Böther, Otto Kißig, Martin Taraz, Sarel Cohen, Karen Seidel, and Tobias Friedrich. What’s wrong with deep learning in tree search for combinatorial optimization. In *International Conference on Learning Representations*, 2021.
- [Burer *et al.*, 2002] Samuel Burer, Renato DC Monteiro, and Yin Zhang. Rank-two relaxation heuristics for max-cut and other binary quadratic programs. *SIAM Journal on Optimization*, 12(2):503–521, 2002.
- [Cappart *et al.*, 2023] Quentin Cappart, Didier Chételat, Elias B Khalil, Andrea Lodi, Christopher Morris, and Petar Velickovic. Combinatorial optimization and reasoning with graph neural networks. *J. Mach. Learn. Res.*, 24:130–1, 2023.
- [Chen and Tian, 2019] Xinyun Chen and Yuandong Tian. Learning to perform local rewriting for combinatorial optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- [Choo *et al.*, 2022] Jinho Choo, Yeong-Dae Kwon, Jihoon Kim, Jeongwoo Jae, André Hottung, Kevin Tierney, and Youngjune Gwon. Simulation-guided beam search for neural combinatorial optimization. *Advances in Neural Information Processing Systems*, 35:8760–8772, 2022.
- [Dorigo *et al.*, 2006] Marco Dorigo, Mauro Birattari, and Thomas Stutzle. Ant colony optimization. *IEEE computational intelligence magazine*, 1(4):28–39, 2006.
- [Dunning *et al.*, 2018] Iain Dunning, Swati Gupta, and John Silberholz. What works best when? a systematic evaluation of heuristics for max-cut and qubo. *INFORMS Journal on Computing*, 30(3):608–624, 2018.
- [Dwivedi and Bresson, 2020] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*, 2020.
- [Erdős *et al.*, 1960] Paul Erdős, Alfréd Rényi, et al. On the evolution of random graphs. *Publ. math. inst. hung. acad. sci.*, 5(1):17–60, 1960.
- [Gagrani *et al.*, 2022] Mukul Gagrani, Corrado Rainone, Yang Yang, Harris Teague, Wonseok Jeon, Roberto Bondesan, Herke van Hoof, Christopher Lott, Weiliang Zeng, and Piero Zappi. Neural topological ordering for computation graphs. *Advances in Neural Information Processing Systems*, 35:17327–17339, 2022.
- [Garmendia *et al.*, 2023] Andoni I Garmendia, Josu Ceberio, and Alexander Mendiburu. Neural improvement heuristics for graph combinatorial optimization problems. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [Glover and Taillard, 1993] Fred Glover and Eric Taillard. A user’s guide to tabu search. *Annals of operations research*, 41(1):1–28, 1993.
- [Grinsztajn *et al.*, 2024] Nathan Grinsztajn, Daniel Furelos-Blanco, Shikha Surana, Clément Bonnet, and Tom Barrett. Winner takes it all: Training performant rl populations for combinatorial optimization. *Advances in Neural Information Processing Systems*, 36, 2024.
- [Gurobi Optimization, LLC, 2023] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.
- [Hottung *et al.*, 2021] André Hottung, Yeong-Dae Kwon, and Kevin Tierney. Efficient active search for combinatorial optimization problems. In *International Conference on Learning Representations*, 2021.
- [Kim *et al.*, 2021] Minsu Kim, Jinkyoo Park, et al. Learning collaborative policies to solve np-hard routing problems. *Advances in Neural Information Processing Systems*, 34:10418–10430, 2021.
- [Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, 2016.
- [Kool *et al.*, 2018] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2018.
- [Kwon *et al.*, 2020] Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. Pomo: Policy optimization with multiple optima for reinforcement learning. *Advances in Neural Information Processing Systems*, 33:21188–21198, 2020.
- [Lamm *et al.*, 2016] Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Renato F Werneck. Finding near-optimal independent sets at scale. In *2016 Proceedings of the Eighteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 138–150. SIAM, 2016.
- [Lawler *et al.*, 1980] Eugene L. Lawler, Jan Karel Lenstra, and AHG Rinnooy Kan. Generating all maximal independent sets: Np-hardness and polynomial-time algorithms. *SIAM Journal on Computing*, 9(3):558–565, 1980.



- [Lawler *et al.*, 1986] Eugene L Lawler, Jan Karel Lenstra, AHG Rinnooy Kan, and David Bernard Shmoys. The traveling salesman problem: a guided tour of combinatorial optimization. *The Journal of the Operational Research Society*, 37(5):535, 1986.
- [Lu *et al.*, 2019] Hao Lu, Xingwen Zhang, and Shuang Yang. A learning-based iterative method for solving vehicle routing problems. In *International conference on learning representations*, 2019.
- [Luo *et al.*, 2024] Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial optimization with heavy decoder: Toward large scale generalization. *Advances in Neural Information Processing Systems*, 36, 2024.
- [Ma *et al.*, 2021] Yining Ma, Jingwen Li, Zhiguang Cao, Wen Song, Le Zhang, Zhenghua Chen, and Jing Tang. Learning to iteratively solve routing problems with dual-aspect collaborative transformer. *Advances in Neural Information Processing Systems*, 34:11096–11107, 2021.
- [Ma *et al.*, 2023] Yining Ma, Zhiguang Cao, and Yeow Meng Chee. Learning to search feasible and infeasible regions of routing problems with flexible neural k-opt. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [Mazyavkina *et al.*, 2021] Nina Mazyavkina, Sergey Sviridov, Sergei Ivanov, and Evgeny Burnaev. Reinforcement learning for combinatorial optimization: A survey. *Computers & Operations Research*, 134:105400, 2021.
- [Mirhoseini *et al.*, 2021] Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Azade Nazi, et al. A graph placement methodology for fast chip design. *Nature*, 594(7862):207–212, 2021.
- [Papadimitriou, 1992] Christos H Papadimitriou. The complexity of the lin–kernighan heuristic for the traveling salesman problem. *SIAM Journal on Computing*, 21(3):450–465, 1992.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [Vrček *et al.*, 2022] Lovro Vrček, Xavier Bresson, Thomas Laurent, Martin Schmitz, and Mile Šikić. Learning to untangle genome assembly with graph convolutional networks. *arXiv preprint arXiv:2206.00668*, 2022.
- [Wang and Tang, 2021] Qi Wang and Chunlei Tang. Deep reinforcement learning for transportation network combinatorial optimization: A survey. *Knowledge-Based Systems*, 233:107526, 2021.
- [Williams, 1992] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992.
- [Wu *et al.*, 2021] Yaoxin Wu, Wen Song, Zhiguang Cao, Jie Zhang, and Andrew Lim. Learning improvement heuristics for solving routing problems. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [Xu and Li, 2000] Ke Xu and Wei Li. Exact phase transitions in random constraint satisfaction problems. *Journal of Artificial Intelligence Research*, 12:93–103, 2000.
- [Ye *et al.*, 2023] Haoran Ye, Jiarui Wang, Zhiguang Cao, Helan Liang, and Yong Li. Deepaco: Neural-enhanced ant systems for combinatorial optimization. *arXiv preprint arXiv:2309.14032*, 2023.
- [Zhang *et al.*, 2023] Dinghuai Zhang, Hanjun Dai, Nikolay Malkin, Aaron Courville, Yoshua Bengio, and Ling Pan. Let the flows tell: Solving graph combinatorial problems with gflownets. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.