

Optimal Solving of Constrained Path-Planning Problems with Graph Convolutional Networks and Optimized Tree Search

Kevin Osanlou^{1,2} Andrei Bursuc³ Christophe Guettier¹ Tristan Cazenave² Eric Jacopin⁴
¹Safran Electronics & Defense ²LAMSADE, Paris-Dauphine University ³valeo.ai
⁴CREC Saint-Cyr, Coetquidan School Campus
{kevin.osanlou, christophe.guettier}@safrangroup.com andrei.bursuc@valeo.com
tristan.cazenave@lamsade.dauphine.fr eric.jacopin@st-cyr.terre-net.defense.gouv.fr

Abstract—Deep learning-based methods are growing prominence for planning purposes. In this paper, we present a hybrid planner that combines a graph machine learning model and an optimal solver based on branch and bound tree search for path-planning tasks. More specifically, a graph neural network is used to assist the branch and bound algorithm in handling constraints associated with a desired solution path. There are multiple downstream practical applications, such as Autonomous Unmanned Ground Vehicles (AUGV), typically deployed in disaster relief or search and rescue operations. In off-road environments, AUGVs must dynamically optimize a source-destination path under various operational constraints, out of which several are difficult to predict in advance and need to be addressed online. We conduct experiments on realistic scenarios and show that graph neural network support enables substantial speedup and smoother scaling to harder path-planning problems. Additionally, information provided by the graph neural network enables the approach to outperform problem-specific handcrafted heuristics, highlighting the potential graph neural networks hold for path-planning tasks.

I. INTRODUCTION

Automated path-planning is an area of interest in AI with a wide panel of applications. The ability to efficiently plan an optimal path in a geometric graph that meets a set of requirements is becoming increasingly critical in a world where autonomy is starting to prevail. The requirements usually consist in a set of constraints imposed on the solution path, making it more difficult to compute. In the case of autonomous unmanned ground vehicles (AUGV), terrain structure is represented through a geometric graph, and maneuvers must consider terrain knowledge. Disaster relief, logistics, or area surveillance are a few among many applications for which online constrained path-planning algorithms enable autonomous mobility using perception and control functionalities. The ability of the AUGV to efficiently come up with an optimal path for a given mission has a direct impact on operational efficiency, underpinning the importance of an efficient path-planner.

For such problems, classical robotic systems integrate A* algorithms [1] as a best-first search approach in the space of available paths. For a complete overview of static algorithms (such as A*), replanning algorithms (D*), anytime algorithms (*e.g.* ARA*), and anytime replanning algorithms (AD*), we refer the reader to Ferguson *et al.* [2]. Tree search algorithms stemming from A* require the specification of a planning domain where constraints are modeled, and can be

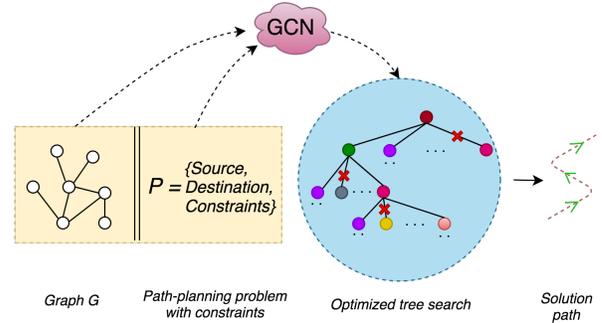


Fig. 1. Proposed framework for solving a path-planning problem with constraints in a graph with a GCN-assisted solver. The GCN takes as input the graph and the problem, and provides relevant information to speed up a tree search, after which an optimal solution path can be built.

more or less efficient depending on the affinity of the search heuristic with the planning context. In this work, we focus on the branch and bound (B&B) tree search algorithm instead [3]. More specifically, our study focuses on the performance gain when it is coupled with machine learning techniques.

Convolutional Neural Networks (CNNs) have been proven to be very efficient for computer vision applications, such as image recognition [4]. They are a multilayer perceptron variant designed for minimal preprocessing and capable of detecting complex patterns in images. In this paper, we are dealing with graphs which represent either maneuvers or off-road navigation. Instead of CNNs, the paper focuses on graph convolutional neural networks (GCNs), a recent architecture for learning complex patterns in graph data [5], [6], [7]. One of the main reasons GCNs are preferred over CNNs for graph processing is that nodes and edges are attributed relevant features. CNNs are not able to learn features of an equivalent quality from an image of the graph. Moreover, CNNs are not invariant to node permutations, an issue GCNs do not share.

In this paper, we study a GCN-based approach for constrained path-planning and stick to exact resolution methods for path-related problems in a specific graph. We run experiments on realistic AUGV scenarios for which we consider mandatory pass-by nodes as type of constraint. This makes path-planning similar to the traveling salesman problem (TSP). The TSP is an NP-hard problem for which there exist efficient approximate solvers [8], however it remains a challenge for exact approaches ([9], [10], [11]). We make the following contributions. First, we define a GCN architecture

suitable for the considered problem. Second, we propose a self-supervised training strategy for the GCN. We then provide a framework which combines the GCN with the depth-first branch and bound (B&B) algorithm. Finally, we conduct experiments on realistic problems for which results exhibit accelerated solving performance.

II. RELATED WORK

In the past few years there has been a growing interest for transferring the intuitions and practices from neural networks on structured inputs towards graphs [12], [13], [14], [6]. [13] bridges spectral graph theory with multi-layer neural networks by learning smooth spectral multipliers of the graph Laplacian. Then [14] and [6] approximate these smooth filters in the spectral domain using polynomials of the graph Laplacian. Free parameters of the polynomials are learned by a neural network, avoiding the costly computation of the eigenvectors of the graph Laplacian. We refer the reader to [15] for a comprehensive review on learning on graphs.

Applications of these types of networks are starting to emerge. Recent works suggest GCNs are capable of making key decisions to solve either path-planning problems in graphs [16], or even STRIPS planning tasks [17]. Regarding path-related optimization problems, *Liet et al.* [16] tackle the maximum independent set (MIS) problem with a solver that combines a GCN-guided tree search and local search optimization. The input to their solver is a graph, and the output is a highly optimized solution. *Kool et al.* [18] propose a reinforcement learning framework to solve the travelling salesman problem (TSP) and variants of the vehicle routing problem (VRP) approximately. While they prefer an encoder-decoder architecture over GCNs, they achieve better results than previous learning-based approaches. These related works focus on the approximate solving of a given task in big graphs. To this end, a learning model is coupled with tree search algorithms to narrow a wide search space in order to retrieve a high quality solution in a short time. In contrast, our work is intended for optimal task solving in smaller graphs. As optimal solving requires visiting most of the search space to ensure proof of optimality, previous approaches are not suitable and we proceed differently.

III. CONTEXT AND PROBLEM FORMALIZATION

We consider a weighted connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, A)$, where \mathcal{V} is the set of vertices, \mathcal{E} the set of edges and A the adjacency matrix of the graph. While in this work we choose to deal with realistic AUGV scenario graphs which are undirected (§VII), our approach is equally applicable to directed graphs. In a typical crisis scenario, the AUGV has to proceed from an initial point to several areas for information gathering, before making its way to a final destination to share its assessment of the ongoing situation. Another frequent scenario consists in delivering food and first-aid equipment to areas likely to be undergoing a shortage. Such

problems can be formalized mathematically. Let I be a path-planning problem instance, defined as follows:

$$I = (start, dest, C)$$

- $start \in \mathcal{V}$ is the index of the start node in \mathcal{G} ,
- $dest \in \mathcal{V}$ is the index of the destination node in \mathcal{G} ,
- C is a set of constraints that need to be satisfied.

Solving I optimally means finding a path p , *i.e.* a sequence of nodes (or edges), which begins from $start$, ends in $dest$, satisfies all constraints in C and minimizes the total weight of the edges included in p . We can consider various types of constraints in C . In this work, we experiment with constraints related to mandatory nodes, which require the solution path to include a given set of nodes $M \in \mathcal{V}$. In the next sections, we will refer to a path-planning problem instance $(start, dest, M)$ simply as an instance. A valid solution path is required to include every node in M at least once. Since the order of visit is not imposed for M , this problem can be assimilated to a TSP variant.

IV. PATH-BUILDING WITH GRAPH CONVOLUTIONAL NETWORKS

In this section, we present our approach for training a graph neural network on a particular graph. We aim to leverage the learning capacity of the graph neural network to approximate the behavior of a model-based planner on the graph.

A. Neural Networks

Neural Networks (NNs) enable multiple levels of abstraction of data by using models with trainable parameters coupled with non-linear transformations of the input data.

In spite of the complex structure of a NN, the main mechanism is straightforward. A *feedforward neural network*, or *multi-layer perceptron (MLP)*, with L layers describes a function $f(\mathbf{x}; \boldsymbol{\theta}) : \mathbb{R}^{d_x} \mapsto \mathbb{R}^{d_y}$ that maps an input vector $\mathbf{x} \in \mathbb{R}^{d_x}$ to an output vector $\mathbf{y} \in \mathbb{R}^{d_y}$. Vector \mathbf{x} is the input data that we need to analyze (*e.g.* an image, a signal, a graph, etc.), while \mathbf{y} is the expected decision from the NN (*e.g.* a class index, a heatmap, etc.). The function f performs L successive operations over the input \mathbf{x} :

$$h^{(l)} = f^{(l)}(h^{(l-1)}; \theta^{(l)}) = \sigma\left(\theta^{(l)}h^{(l-1)} + b^{(l)}\right) \quad (1)$$

where $h^{(l)}$ is the hidden state of the network and $f^{(l)}$ is the mapping function performed at layer l and parameterized by trainable parameters $\theta^{(l)}$ and bias $b^{(l)}$, and piece-wise activation function $\sigma(\cdot)$; $h^{(0)} = \mathbf{x}$.

CNNs [19], [20] are a popular architecture for 2D data. They generalize MLPs by sliding groups of parameters across an input vector similarly to filters in image processing, leveraging fewer parameters and parallel computation. Hidden states in CNNs preserve the number of dimensions of the input, *i.e.* 2D when images are used as input, and are called *feature maps*.

B. Graph Convolutional Networks

GCNs are generalizations of CNNs to non-Euclidean graphs [15]. GCNs are in fact neural networks based on local operators on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ which are derived from spectral graph theory. The filter parameters are typically shared over all locations in the graph, thus the name *convolutional*.

We consider here the approach of Kipf and Welling [6]. The GCNs have the following layer propagation rule:

$$h^{(l+1)} = \sigma\left(\tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}h^{(l)}\theta^{(l)}\right), \quad (2)$$

where $\tilde{A} = A + I_N$ is the adjacency matrix of the graph with added self-connections such that when multiplying with \tilde{A} we aggregate features vectors from both a node and its neighbors. Matrix I_N is the identity matrix; \tilde{D} is the diagonal node degree matrix of \tilde{A} ; $\sigma(\cdot)$ is the activation function, which we set to $\text{ReLU}(\cdot) = \max(0, \cdot)$. Matrix \tilde{D} is employed for normalization of \tilde{A} in order to avoid a change of scales in the feature vectors when multiplying with \tilde{A} . In [6], the authors argue that using a symmetric normalization, *i.e.* $\tilde{D}^{-\frac{1}{2}}\tilde{A}\tilde{D}^{-\frac{1}{2}}$, ensures better dynamics compared to simple averaging of neighboring nodes in the one-sided normalization $\tilde{D}^{-1}\tilde{A}$.

C. Problem Instance Encoding

The input of our model is a vector \mathbf{x} containing information about the problem instance, including the graph representation. For every instance $I = (start, dest, M)$ of a given graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we associate a vector \mathbf{x} made up of triplet features from each node in \mathcal{G} , making up for a total of $3 \times |\mathcal{V}|$ features. The three features for a node $j \in \mathcal{V}$ are:

- *start node* feature $st_j = 1$ if node j is the start node in instance I , otherwise 0
- *end node* feature $ed_j = 1$ if node j is the end node in instance I , otherwise 0
- *mandatory node* feature $my_j = 1$ if node j is a mandatory node in instance I , otherwise 0.

We obtain the input $\mathbf{x} \in \mathbb{R}^{|\mathcal{V}| \times 3}$ by stacking node features:

$$\mathbf{x} = [[st_1, ed_1, my_1], [st_2, ed_2, my_2], \dots, [st_{|\mathcal{V}|}, ed_{|\mathcal{V}|}, my_{|\mathcal{V}|}]] \quad (3)$$

D. Neural network architecture and training

We define a neural network f that consists of a sequence of multiple graph convolutions followed by a fully connected layer. This GCN takes as input any instance I over the graph \mathcal{G} , and outputs a probability vector $\hat{\mathbf{y}} \in \mathbb{R}^{|\mathcal{V}|}$. Then, $\arg \max(\hat{\mathbf{y}})$ corresponds to the next mandatory node to visit from the *start* node in an optimal path that solves I . The hidden states of the graph convolution layers $h^{(l)} \in \mathbb{R}^{|\mathcal{V}| \times N_{hidden}}$ consist of N_{hidden} higher-dimensional features for each node in the graph. After the final convolutional layer, the feature matrix is flattened into a vector by concatenating its rows and linked to a fully connected layer that maps it to a vector $z \in \mathbb{R}^{|\mathcal{V}|}$. We use the *softmax* function to convert z into probabilities.

NNs trained in a supervised manner use labeled training data, *i.e.* a set of input-output pairs $(\mathbf{x}_i, \mathbf{y}_i)$ sampled from

a large training set. Here, \mathbf{x}_i is an instance I and \mathbf{y}_i is the index of the next mandatory node for \mathbf{x}_i in an optimal path solving I . We train the GCN on instances that have already been optimally solved by an exact planner, which serves here as a *teacher*. The network learns to approximate the solutions computed by the planner. To this end, we train the network using the *negative log-likelihood loss*, used for multi-class classification, and stochastic gradient descent (SGD).

E. Mandatory Node Ordering

For a given input instance, the GCN computes a mandatory node prediction at a time. In order to make it compatible with instances with varying amounts of mandatory nodes we make a few adjustments. Given an instance with multiple mandatory nodes $(start, dest, M)$, we perform multiple recursive GCN calls to get the next mandatory node predictions q_i . More specifically, after a prediction q_i is computed, we generate a sub-instance where the start node becomes the current predicted mandatory node q_i and where the list of mandatory nodes contains the remaining nodes except q_i , *i.e.* $(q_i, dest, \tilde{M}_i = \tilde{M}_{i-1} \setminus q_i)$, where $\tilde{M}_0 = M$ and $q_0 = start$. We use the recursive calls during both training and testing. An interesting side-effect of this strategy is that it ensures an implicit balancing of the training samples by difficulty, as we generate sub-instances ranging from challenging (large $|M|$) to trivial (small $|M|$).

V. SELF-SUPERVISED LEARNING

We present a self-supervised learning strategy aimed at training the GCN for a particular graph. First, we define a planning domain to solve instances. Second, we introduce a modified version of the A* algorithm for generating optimal data on which the GCN is trained. In this section, we refer to an instance as a planning state $s = (start, dest, M)$. An end state is a termination instance, *i.e.* an instance for which the destination node has been reached and all mandatory nodes have been visited. We denote the termination instances as $F_i = (i, i, \emptyset)$, $i \in \{1, 2, \dots, |V|\}$. There are exactly as many termination instances as there are nodes in \mathcal{G} . We respectively define the *successors* and *predecessors* of a state $s = (start, dest, M)$ as $succ(s)$ and $pred(s)$ in Table I.

TABLE I
TRANSITION RULES TO SUCCESSORS AND PREDECESSORS.

$s = (start, dest, M)$	
Successor state s'	Predecessor state s'
$s' = (start', dest', M')$	$s' = (start', dest', M')$
$(start, start') \in \mathcal{E}$	$(start', start) \in \mathcal{E}$
$dest' = dest$	$dest' = dest$
$M' = M \setminus \{start'\}$	$M' = M$ \triangleright ^{1*}
-	or
-	$M' = M \cup \{start\}$ \triangleright ^{2*}
s is not an end state	s' is not an end state
Transition cost: $(start, start')$	Transition cost: $(start', start)$

^{1*} only if $start' \notin M$

^{2*} only if $start \neq dest$

Algorithm 1 Backwards A* for reverse instance solving

```
1: function ComputePaths()
2:   while elapsed time < timeout do
3:     remove state  $s$  from the front of OPEN;
4:     if  $\text{length}(\text{mand}(s)) > 0$  then ▷ 3*
5:       insert the pair  $\langle s, d(s) \rangle$  to the training set
6:       for all  $s' \in \text{pred}(s)$  do ▷ 3*
7:         if  $g(s') > g(s) + c(s, s')$  then
8:            $g(s') = g(s) + c(s, s')$ 
9:            $a(s') = s$ 
10:        if  $\text{mand}(s) = \text{mand}(s')$  then
11:           $d(s') = d(s)$  ▷ 4*
12:        else
13:           $d(s') = \text{startNode}(s)$  ▷ 3,5*
14:        insert  $s'$  into OPEN with value  $g(s')$  ▷ 6*
15: function main()
16:   for all  $i \in \{1, 2, \dots, |V|\}$  do
17:     for all  $s \in S$  do
18:        $g(s) = \infty$ 
19:        $g(E_i) = 0, a(E_i) = \emptyset, d(E_i) = \emptyset$ 
20:        $\text{OPEN} = \emptyset$ 
21:       insert  $F_i$  into OPEN with value  $g(E_i)$  ▷ 6*
22:       ComputePaths();
```

^{3*} $\text{length}(x)$: returns length of list x ; $\text{mand}(x)$: returns mandatory nodes of instance x ; $\text{startNode}(x)$: returns the start node of instance x ; $\text{pred}(x)$: returns all possible predecessors of instance x

^{4*} s' is created from s without adding its start node to the list of mandatory nodes

^{5*} s' is created from s by adding its start node to the list of mandatory nodes

^{6*} $h = 0$

Transition costs from a state to a neighboring state is the cost of the edge in the graph linking the start nodes of both states. With these rules, the destination node remains always the same. Therefore, we run the backwards version of A* from every termination instance F_i as initial state (using pred as rule of succession). For each state s visited by A*, a path p is built from s to F_i which the algorithm considers the shortest. We define $g(s)$ as the cost of p , $a(s)$ as the next state visited after s in p , and $d(s)$ as the first mandatory node of s that is visited in p .

Furthermore we perform the following changes to A*. First, when a shorter path is found to a state s' while developing a state s , *i.e.* $g(s') > g(s) + c(s, s')$, the values of $a(s')$ and $d(s')$ are also updated along with $g(s')$ to take into account the shorter path. Secondly, we set the heuristic function h to 0. Since A* is run backwards from a termination state F_i , we are not aiming for the algorithm to reach a defined state in particular, but seek to reach as many states as possible. This ensures that when a state s is taken from the OPEN priority list of states left to develop, an optimal path from s to F_i has already been found. Choosing $d(s)$ as the next mandatory node to visit thus enables optimal solving. Consequently we can add the pair $\langle s, d(s) \rangle$ to the training set. We provide the pseudo-code in Algorithm 1.

The data generated by the algorithm is added to the training set and shuffled. The GCN is then trained on this set with supervised learning. Results show that training on

the "synthetic" data generated with A* enables the GCN to generalize well on instances that A* did not process. We argue this is because the distribution obtained with A* is related to the path length of resolved instances. In fact, graph patterns already explored in short path solutions are incrementally included into longest ones.

VI. DEPTH-FIRST BRANCH AND BOUND TREE SEARCH

It is possible to resolve an instance within a combinatorial search tree, rather than in the path-planning domain defined in §V. To this end, we compute the shortest source-destination paths for every pair of nodes $\langle i, j \rangle$ in \mathcal{G} using Dijkstra's algorithm, as well as the associated path cost. Solving an instance $I = (\text{start}, \text{dest}, M)$ then becomes equivalent to finding the optimal order in which the mandatory nodes in M are visited for the first time. The solution path associated with an order $o = (m_1, m_2, m_3, \dots, m_{q-1}, m_q)$ of the mandatory nodes can be built by concatenating the shortest path from start to m_1 , from m_1 to m_2 , from m_2 to m_3 , ... , from m_{q-1} to m_q , and from m_q to dest . Its total cost is the sum of the cost of each shortest path used to build it. A particular tree can be searched for identifying an optimal order of the mandatory nodes. For an instance I , we define the root of this tree as the start node, every leaf node as the dest node, and every intermediate tree node as a mandatory node in M , such that a path from the root of the tree to a leaf defines an order in which to visit the mandatory nodes in M . The cost of transitioning from a node v to a child node v' in the tree is the cost of the shortest path in \mathcal{G} between the pair $\langle v, v' \rangle$. In the following, we refer to this tree as the mandatory search tree. Since we are using here only mandatory node constraints, the combinatorial optimization problem associated with this tree displays a similar structure with the TSP. However, it differs in the choice of the start and destination nodes which are fixed.

The branch and bound algorithm (B&B) is a popular tree search algorithm that is well known for its computational efficiency ([3]). In this work, we consider a depth-first B&B search algorithm. When developing a node inside the tree, the algorithm checks if each branch is expected to host a better solution than the best solution found so far. Should that not be the case for a given branch, the branch is cut, and the algorithm will not develop nodes further down the branch. This is done by using a lower bound and an upper bound. The lower bound is the sum of the total cost from the root node to the current node and a heuristic function π that approximates the remaining cost from the current node to the best achievable solution in branches below. This heuristic function should return a value as close as possible to this remaining cost (to cut as frequently as possible), while staying smaller (for the algorithm to remain optimal).

Next, we define the heuristic function π that we use. Let $v \in M$ be a mandatory node in the mandatory search tree, and R the set of remaining mandatory nodes left to the leaf node dest , *i.e.* the nodes in M that haven't been included between the root node start and v . Let $D = \{v\} \cup R \cup \{\text{dest}\}$. We define two functions, \min_1 and \min_2 , that respectively

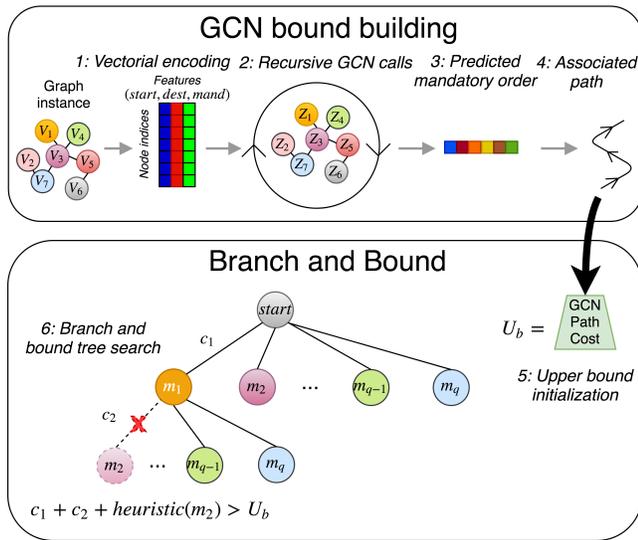


Fig. 2. GCN-assisted branch and bound algorithm pipeline. The GCN is used recursively to build an order of visit for the mandatory nodes. The order is converted into a path using the previously computed shortest path pairs, and the cost of the path is used as an initial upper bound for the algorithm. Here, the upper bound of the GCN allows for a level 2 early cut.

return for a node x in D the lowest shortest path cost in \mathcal{G} from x to any node in $D \setminus \{x\}$, and the second lowest such cost. We build the heuristic function π by considering the remaining nodes left. For each node left, we consider the weight of all edges connecting it to other nodes in D and add the first and second smallest such weights, with the exception of v and $dest$ for which only the first smallest weight is added, and divide the total by 2:

$$\pi(v) = \frac{1}{2}(\min_1(v) + \min_1(dest) + \sum_{r \in R} \min_1(r) + \min_2(r))$$

Recent progress in learning-assisted tree search has shown that machine learning can be used to narrow the search space in very large domains to allow for efficient solving. Inspiring results have been shown in the game of Go [21], [22], where a very good solution, which is not required to be optimal, is found in record time. Promising parts of the search tree are first visited in accordance with the suggestions of the neural network, unpromising parts then, only if time allows. On the other hand, in a context where finding an optimal solution is critical, the search cannot be directed in such a way, as proof of optimality is required. Consequently, we keep the GCN out of the tree search procedure. For a given instance, we use the GCN recursively in order to obtain a suggested order of visit of the mandatory nodes (IV-E), from which we build the associated solution path by concatenating the shortest paths. This is done in negligible time compared to the tree search. The cost of the solution path found in such *probing* manner [23] is then used as an initial upper bound for the B&B algorithm (figure 2). In section §VII we conduct experiments to evaluate the influence of the upper bound obtained with the GCN on search performance.

A. Benchmarks and baselines

We run experiments to evaluate the impact of the GCN’s upper bound on the B&B method. Since proof of optimality is necessary in our context, we focus only on small-scale problems for which optimal solving is possible in reasonable time. We consider four different graphs, \mathcal{G}_1 , \mathcal{G}_2 , \mathcal{G}_3 and \mathcal{G}_4 , with respectively 15, 23, 22 and 23 nodes. These graphs represent realistic AUGV crisis scenarios in which aid has to be provided to key points in operational areas. More details on the graphs are available in [24], from which the scenarios have been built. We generate 1512, 2928, 2712 and 2712 random instances for each graph respectively. In order to remain close to some ‘realistic’ instances, we generate the instances as follows: using the shortest source-destination paths computed previously, we apply a decimation ratio (typically 80%) to keep only the 20% source-destination pairs that have the longest shortest paths. For each resulting pair $\langle start, dest \rangle$ kept, we generate multiple random instances with an increasing cardinality for the set of mandatory nodes, ranging from 5 to 12.

We consider 4 different baseline solvers on the benchmark instances generated to compare solving performance. All solvers search for an optimal solution path. First, we use a solver based on dynamic programming (DP) which searches the mandatory search tree. Second, we run the B&B algorithm to search the mandatory search tree, both with and without the upper bound provided by the GCN. Lastly, we solve instances using forward A* applied on the planning domain described in section §V, with the minimum spanning tree (MST) as heuristic function. The MST heuristic is computed for an instance $(start, dest, M)$ by considering the complete graph \mathcal{G}' , which comprises only the $start$, $dest$ and mandatory nodes M . All pairs of nodes (v, v') in \mathcal{G}' are connected by an edge which has a weight equal to the cost of the shortest path from v to v' in \mathcal{G} . The MST heuristic value is obtained by adding the following three values: the total weight of the MST of all mandatory nodes M in \mathcal{G}' , the minimum edge weight in \mathcal{G}' from the $start$ node to any node in the MST, and the minimum edge weight in \mathcal{G}' from any node in the MST to the $dest$ node.

B. Implementation details

We set the run-time of A* to 10 hours per graph for data generation. We use 3 graph convolutional layers of width 100. During training, we apply batch normalization [25] with decay of moving average $\varepsilon = 0.9$, dropout with *drop rate* of 0.1, and train the GCN with *Adam* [26]. We set the learning rate to $\eta = 10^{-4}$. We train models on a Tesla P100 GPU using over 1.5M training examples generated by A*, for ~ 5 hours. We conduct the benchmark tests in this section on a laptop with an Intel i5 processor and 8GB of RAM. We point out that our approach requires GPU only for the training of the GCN, which can be done offline. Problem instances can be then solved online on a CPU.

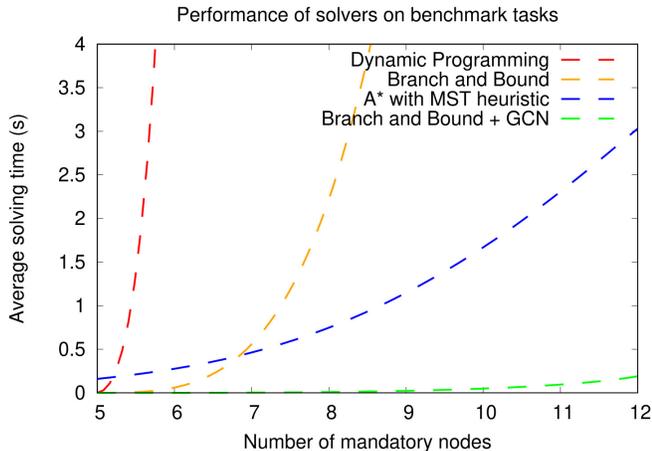


Fig. 3. Comparison of the performance of different solvers on benchmark instances generated for graph \mathcal{G}_2 . The X axis represents the number of mandatory nodes of the instances, the Y axis the average solving time. We limit the Y axis to a range [0,4] to obtain a better comparison scale.

C. Results

We summarize results for \mathcal{G}_2 in Figure 3. We detail the experiments for graph \mathcal{G}_3 and \mathcal{G}_4 in Tables II and III. We include solving time only if it is measurable by CPU clock time. Figures for all graphs show a similar trend. We note that for instances with seven mandatory nodes and more, best-first algorithms such as A* applied on the planning domain defined in §V become more suited than depth-first algorithms such as B&B applied on the combinatorial search tree associated with the mandatory constraints. Although each planning state takes longer to compute in order to account for the specifics of the planning domain, overall significantly fewer planning states are visited than mandatory search tree nodes. This is because the planning domain takes advantage of the graph structure, which gives A* a significant edge over depth-first DP and B&B. We note, however, that when the upper bound of the GCN is used, the B&B algorithm is able to outperform A* on all instances, even the most complex ones, while scaling more smoothly with the number of mandatory nodes. In table IV we provide additional

TABLE II

EXPERIMENTS FOR GRAPH \mathcal{G}_3 . LEGEND: T/O = 5 MINUTES TIMEOUT.

Mandatory #:	5	7	9	11	12
DP					
Avg. node visits:	11.7K	876K	98.6M	T/O	T/O
Avg. time (s):	-	0.26	30.52	T/O	T/O
B&B					
Avg. node visits:	418	4,94K	146K	11,1M	120M
Avg. time (s):	-	-	0.07	5.39	71.1
A*, h=MST					
Avg. state visits:	26.4	58.5	141	342	503
Avg. time (s):	0.14	0.29	0.73	1.84	2.70
B&B + GCN					
Avg. node visits:	148	1,24K	10,8K	161K	642K
Avg. time (s):	-	-	0.01	0.08	0.34

TABLE III

EXPERIMENTS FOR GRAPH \mathcal{G}_4 . LEGEND: T/O = 5 MINUTES TIMEOUT.

Mandatory #:	5	7	9	11	12
DP					
Avg. node visits:	11.7K	876K	98.6M	T/O	T/O
Avg. time (s):	-	0.24	30.53	T/O	T/O
B&B					
Avg. node visits:	545	7.58K	193K	11,8M	122M
Avg. time (s):	-	-	0.09	5.28	70.8
A*, h=MST					
Avg. state visits:	41.78	96.8	226	787	1,09K
Avg. time (s):	0.15	0.44	0.96	3.18	4.36
B&B + GCN					
Avg. node visits:	233	2.46K	29.1K	235K	859K
Avg. time (s):	-	-	0.02	0.11	0.56

insight on these results through information collected from the mandatory search tree for instances with 11 mandatory nodes. The average number of nodes processed in the subtree of each child node of the root node is given, as well as the average score of the best known solution after the subtree is processed.

TABLE IV

INFORMATION FROM THE MANDATORY SEARCH TREE.

Root node child #	-	1	2	3	...
B&B					
Avg. node visits	-	11M	50K	30K	...
Avg. best sol. score	∞	10.35K	9940	9781	...
B&B + GCN					
Avg. node visits	-	74K	11K	7K	...
Avg. best sol. score	9890	9300	9264	9249	...

Since our B&B algorithm is depth-first, processing the entire subtree under the first child node of the root node when no initial upper bound is known is highly computationally expensive. Indeed, no cut can be made until a leaf node is reached, and even then, the identified solution is very likely to be costly compared to the optimal solution, thus the updated upper bound would still not allow for frequent cuts, until a good part of the subtree has been processed. On the other hand, if a good upper bound is known in advance, which is generally the case for the one given by the GCN in our experiments, the algorithm does not suffer from this issue, and early cuts can be made.

VIII. DISCUSSION AND FURTHER WORKS

We experiment with path-planning problems defined by three features: the start node, destination node, and mandatory nodes. We accelerate optimal depth-first solving of the search tree associated with the mandatory constraints by leveraging the upper bound computed by the GCN. We show that this speedup is significant, competing successfully with A*. This is the case even for scenarios where handling constraints within the planning domain is more appropriate than extracting and solving them separately. Also, our attempts to guide A* search with a GCN heuristic achieved worse results

than the MST heuristic. The reason is due to the best-first approach for which the GCN is unable to provide a suitable heuristic. Moreover, the proposed framework can include additional types of constraints for path-planning problems. Each new constraint type results in additional features on nodes, and potentially also on edges [27]. In this case, the GCN would learn to predict the next node to visit, and not the next mandatory node. Recursive GCN calls would be made until a solution which satisfies all constraints is found, backtracking when necessary. The approach can be combined with state-of-the-art constraint propagation techniques [28]. In the same manner, the solution cost can be used as an initial upper bound for a depth-first search of a combinatorial tree associated with the constraints.

Learning-wise, the more constraint types there are for a path-planning problem, the wider the GCN learning domain will be. Further work will especially focus on this limitation to relate the exhaustiveness of the training phase with the variety of constraint types. Also, the proposed approach requires neural network offline training for a given graph (e.g. problem scenario). It can then be used online for path re-planning purposes, as the AUGV drives through the graph, with appealing computational performances.

IX. CONCLUSION

We introduced a method combining graph neural networks and branch and bound (B&B) tree search to handle constraints in path-planning, successfully accelerating optimal solving of path-planning tasks. A relevant self-supervised strategy has been developed, based on A*, which provides appropriate data to train the graph neural network. The heuristic information computed by the graph neural network enables better scaling of the B&B algorithm onto more complex problems. Results exhibit solving times that outperform A* with problem-specific handcrafted heuristics. Various path-planning applications for autonomous vehicles can benefit from such an approach, especially when known terrains are given and path or itineraries must be computed on the fly. We also hope this line of work will serve to highlight the merits of using graph neural networks for path-planning tasks.

REFERENCES

- [1] P. Hart, N. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems, Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [2] D. Ferguson, M. Likhachev, and A. T. Stentz, "A guide to heuristic-based path planning," in *Proceedings of the International Workshop on Planning under Uncertainty for Autonomous Systems, International Conference on Automated Planning and Scheduling (ICAPS)*, June 2005.
- [3] P. M. Narendra and K. Fukunaga, "A branch and bound algorithm for feature subset selection," *IEEE Transactions on Computers*, vol. C-26, pp. 917–922, 1977.
- [4] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [5] J. Bruna, W. Zaremba, A. Szlam, and Y. Lecun, "Spectral networks and locally connected networks on graphs," in *International Conference on Learning Representations*, 2014.
- [6] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations*, 2017.
- [7] Y. Li, D. Tarlow, M. Brockschmidt, and R. S. Zemel, "Gated graph sequence neural networks," in *International Conference on Learning Representations*, 2016.
- [8] J. Monnot, V. T. Paschos, and S. Toulouse, "Approximation algorithms for the traveling salesman problem," *Mathematical methods of operations research*, vol. 56, no. 3, pp. 387–405, 2003.
- [9] Y. Caseau and F. Laburthe, "Solving small ttps with constraints," in *ICLP*, vol. 97, 1997, p. 104.
- [10] T. Volgenant and R. Jonker, "A branch and bound algorithm for the symmetric traveling salesman problem based on the 1-tree relaxation," *European Journal of Operational Research*, vol. 9, pp. 83–89, 01 1982.
- [11] G. Laporte, "The traveling salesman problem: An overview of exact and approximate algorithms," *European Journal of Operational Research*, vol. 59, no. 2, pp. 231–247, 1992.
- [12] M. Gori, G. Monfardini, and F. Scarselli, "A new model for learning in graph domains," in *Neural Networks, 2005. IJCNN'05. Proceedings. 2005 IEEE International Joint Conference on*, vol. 2. IEEE, 2005, pp. 729–734.
- [13] M. Henaff, J. Bruna, and Y. LeCun, "Deep convolutional networks on graph-structured data," *arXiv preprint arXiv:1506.05163*, 2015.
- [14] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in Neural Information Processing Systems*, 2016, pp. 3844–3852.
- [15] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: going beyond euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, 2017.
- [16] Z. Li, Q. Chen, and V. Koltun, "Combinatorial optimization with graph convolutional networks and guided tree search," in *Advances in Neural Information Processing Systems*, 2018, pp. 536–545.
- [17] T. Ma, P. Ferber, S. Huo, J. Chen, and M. Katz, "Adaptive planner scheduling with graph neural networks," *CoRR*, vol. abs/1811.00210, 2018. [Online]. Available: <http://arxiv.org/abs/1811.00210>
- [18] W. Kool and M. Welling, "Attention solves your tsp," *arXiv preprint arXiv:1803.08475*, 2018.
- [19] K. Fukushima and S. Miyake, "Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition," in *Competition and cooperation in neural nets*. Springer, 1982, pp. 267–285.
- [20] Y. LeCun, Y. Bengio, *et al.*, "Convolutional networks for images, speech, and time series," *The handbook of brain theory and neural networks*, vol. 3361, no. 10, p. 1995, 1995.
- [21] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, *et al.*, "Mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [22] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, p. 354, 2017.
- [23] K. Osanlou, C. Guettier, A. Bursuc, T. Cazenave, and E. Jacopin, "Constrained shortest path search with graph convolutional neural networks," in *Workshop on Planning and Learning (PAL-18)*, 2018.
- [24] C. Guettier, "Solving planning and scheduling problems in network based operations," in *Proceedings of Constraint Programming (CP)*, 2007.
- [25] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference for Learning Representations*, 2015.
- [27] W. B. W. Vos, "End-to-end learning of latent edge weights for graph convolutional networks," Master's thesis, University of Amsterdam, 8 2017.
- [28] C. Guettier and F. Lucas, "A constraint-based approach for planning unmanned aerial vehicle activities," *The Knowledge Engineering Review*, vol. 31, no. 5, pp. 486–497, 2016.
- [29] K. Osanlou, C. Guettier, A. Bursuc, T. Cazenave, and E. Jacopin,

- “Learning-based preference prediction for constrained multi-criteria path-planning,” *arXiv preprint arXiv:2108.01080*, 2021.
- [30] G. Li, M. Müller, B. Ghanem, and V. Koltun, “Training graph neural networks with 1000 layers,” in *International conference on machine learning*. PMLR, 2021, pp. 6437–6449.
- [31] Z. Zhao, G. Verma, C. Rao, A. Swami, and S. Segarra, “Distributed scheduling using graph neural networks,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 4720–4724.
- [32] J. Park, J. Chun, S. H. Kim, Y. Kim, and J. Park, “Learning to schedule job-shop problems: representation and policy learning using graph neural network and reinforcement learning,” *International Journal of Production Research*, vol. 59, no. 11, pp. 3360–3377, 2021.
- [33] C. Zhang, W. Song, Z. Cao, J. Zhang, P. S. Tan, and X. Chi, “Learning to dispatch for job shop scheduling via deep reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 1621–1632, 2020.
- [34] Z. Wang and M. Gombolay, “Learning scheduling policies for multi-robot coordination with graph attention networks,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4509–4516, 2020.
- [35] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, “Graph neural networks for decentralized multi-robot path planning,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 785–11 792.
- [36] F. Zhou, Q. Yang, T. Zhong, D. Chen, and N. Zhang, “Variational graph neural networks for road traffic prediction in intelligent transportation systems,” *IEEE Transactions on Industrial Informatics*, vol. 17, no. 4, pp. 2802–2812, 2020.
- [37] T. Silver, R. Chitnis, A. Curtis, J. Tenenbaum, T. Lozano-Perez, and L. P. Kaelbling, “Planning with learned object importance in large problem instances using graph neural networks,” *arXiv preprint arXiv:2009.05613*, 2020.
- [38] M. S. A. Hameed and A. Schwung, “Reinforcement learning on job shop scheduling problems using graph networks,” *arXiv preprint arXiv:2009.03836*, 2020.
- [39] F. Zhou, Q. Yang, K. Zhang, G. Trajcevski, T. Zhong, and A. Khokhar, “Reinforced spatiotemporal attentive graph neural networks for traffic forecasting,” *IEEE Internet of Things Journal*, vol. 7, no. 7, pp. 6414–6428, 2020.
- [40] Q. Cappart, D. Chételat, E. Khalil, A. Lodi, C. Morris, and P. Veličković, “Combinatorial optimization and reasoning with graph neural networks,” *arXiv preprint arXiv:2102.09544*, 2021.
- [41] R. Sato, M. Yamada, and H. Kashima, “Approximation ratios of graph neural networks for combinatorial problems,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [42] M. Prates, P. H. Avelar, H. Lemos, L. C. Lamb, and M. Y. Vardi, “Learning to solve np-complete problems: A graph neural network for decision tsp,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 4731–4738.
- [43] L. Hu, Z. Liu, W. Hu, Y. Wang, J. Tan, and F. Wu, “Petri-net-based dynamic scheduling of flexible manufacturing system via deep reinforcement learning with graph convolutional network,” *Journal of Manufacturing Systems*, vol. 55, pp. 1–14, 2020.
- [44] Q. Li, W. Lin, Z. Liu, and A. Prorok, “Message-aware graph attention networks for large-scale multi-robot path planning,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5533–5540, 2021.
- [45] K. Rusek and P. Cholda, “Message-passing neural networks learn little’s law,” *IEEE Communications Letters*, vol. 23, no. 2, pp. 274–277, 2018.
- [46] Y. Hu, S. Chen, Y. Zhang, and X. Gu, “Collaborative motion prediction via neural motion message passing,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 6319–6328.
- [47] F. Gama, E. Tolstaya, and A. Ribeiro, “Graph neural networks for decentralized controllers,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5260–5264.
- [48] Y. Hu, Y. Yao, and W. S. Lee, “A reinforcement learning approach for optimizing multiple traveling salesman problems over graphs,” *Knowledge-Based Systems*, vol. 204, p. 106244, 2020.
- [49] X. Weng, Y. Yuan, and K. Kitani, “Joint 3d tracking and forecasting with graph neural network and diversity sampling,” *arXiv preprint arXiv:2003.07847*, vol. 2, no. 6.2, p. 1, 2020.
- [50] H. Lemos, M. Prates, P. Avelar, and L. Lamb, “Graph colouring meets deep learning: Effective graph neural network models for combinatorial problems,” in *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 2019, pp. 879–885.
- [51] I. Drori, A. Kharkar, W. R. Sickinger, B. Kates, Q. Ma, S. Ge, E. Dolev, B. Dietrich, D. P. Williamson, and M. Udell, “Learning to solve combinatorial optimization problems on real-world graphs in linear time,” in *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2020, pp. 19–24.
- [52] N. Vesselinova, R. Steinert, D. F. Perez-Ramirez, and M. Boman, “Learning combinatorial optimization on graphs: A survey with applications to networking,” *IEEE Access*, vol. 8, pp. 120 388–120 416, 2020.
- [53] T. Guo, C. Han, S. Tang, and M. Ding, “Solving combinatorial problems with machine learning methods,” *Nonlinear Combinatorial Optimization*, 2019.
- [54] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, “A comprehensive survey on graph neural networks,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, pp. 4–24, 2019.
- [55] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, “How powerful are graph neural networks?” *ArXiv*, vol. abs/1810.00826, 2019.
- [56] M. Liu, H. Gao, and S. Ji, “Towards deeper graph neural networks,” *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020.
- [57] R. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, “Gnnexplainer: Generating explanations for graph neural networks,” *Advances in neural information processing systems*, vol. 32, pp. 9240–9251, 2019.
- [58] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, and Z. Zhang, “Deep graph library: A graph-centric, highly-performant package for graph neural networks,” *arXiv: Learning*, 2019.
- [59] E. Rossi, F. Frasca, B. P. Chamberlain, D. Eynard, M. M. Bronstein, and F. Monti, “Sign: Scalable inception graph neural networks,” *ArXiv*, vol. abs/2004.11198, 2020.
- [60] V. K. Garg, S. Jegelka, and T. Jaakkola, “Generalization and representational limits of graph neural networks,” *ArXiv*, vol. abs/2002.06157, 2020.
- [61] A. Loukas, “What graph neural networks cannot learn: depth vs width,” *ArXiv*, vol. abs/1907.03199, 2020.
- [62] F. Lin and H.-P. Hsieh, “A goal-prioritized algorithm for additional route deployment on existing mass transportation system,” in *2020 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2020, pp. 1130–1135.
- [63] F. Lin, J.-Y. Fang, and H.-P. Hsieh, “A gaussian-prioritized approach for deploying additional route on existing mass transportation with neural-network-based passenger flow inference,” in *2020 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2020, pp. 1–8.
- [64] S. Sombolstan, A. Rasooli, and S. Khodaygan, “Optimal path-planning for mobile robots to find a hidden target in an unknown environment based on machine learning,” *Journal of Ambient Intelligence and Humanized Computing*, vol. 10, no. 5, pp. 1841–1850, 2019.
- [65] M. W. Otte, “A survey of machine learning approaches to robotic path-planning,” *University of Colorado at Boulder, Boulder*, 2015.
- [66] T. Zohdi, “The game of drones: rapid agent-based machine-learning models for multi-uav path planning,” *Computational Mechanics*, vol. 65, no. 1, pp. 217–228, 2020.
- [67] R. Yonetani, T. Taniai, M. Barekatin, M. Nishimura, and A. Kanezaki, “Path planning using neural a* search,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 12 029–12 039.
- [68] I. Sung, B. Choi, and P. Nielsen, “On the training of a neural network for online path planning with offline path planning algorithms,” *International Journal of Information Management*, vol. 57, p. 102142, 2021.
- [69] D. Xin, C. Hua-hua, and G. Wei-kang, “Neural network and genetic algorithm based global path planning in a static environment,” *Journal of Zhejiang University-Science A*, vol. 6, no. 6, pp. 549–554, 2005.
- [70] H. Zhou, D. Ren, H. Xia, M. Fan, X. Yang, and H. Huang, “Ast-gnn: An attention-based spatio-temporal graph neural network for interaction-aware pedestrian trajectory prediction,” *Neurocomputing*, vol. 445, pp. 298–308, 2021.

- [71] K. Rusek, J. Suárez-Varela, P. Almasan, P. Barlet-Ros, and A. Cabellos-Aparicio, "Routenet: Leveraging graph neural networks for network modeling and optimization in sdn," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 10, pp. 2260–2270, 2020.
- [72] Y. Zhang, S. Li, and H. Guo, "A type of biased consensus-based distributed neural network for path planning," *Nonlinear Dynamics*, vol. 89, no. 3, pp. 1803–1815, 2017.
- [73] R. Glasius, A. Komoda, and S. C. Gielen, "Neural network dynamics for path planning and obstacle avoidance," *Neural Networks*, vol. 8, no. 1, pp. 125–133, 1995.
- [74] N. Wu, X. W. Zhao, J. Wang, and D. Pan, "Learning effective road network representation with hierarchical graph neural networks," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 6–14.
- [75] Z. Li, C. Lu, Y. Yi, and J. Gong, "A hierarchical framework for interactive behaviour prediction of heterogeneous traffic participants based on graph neural network," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [76] S. Chen, J. Dong, P. Ha, Y. Li, and S. Labi, "Graph neural network and reinforcement learning for multi-agent cooperative control of connected autonomous vehicles," *Computer-Aided Civil and Infrastructure Engineering*, vol. 36, no. 7, pp. 838–857, 2021.
- [77] D. Cao, J. Li, H. Ma, and M. Tomizuka, "Spectral temporal graph neural network for trajectory prediction," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1839–1845.
- [78] U. A. Syed, F. Kunwar, and M. Iqbal, "Guided autowave pulse coupled neural network (gapcnn) based real time path planning and an obstacle avoidance scheme for mobile robots," *Robotics and autonomous systems*, vol. 62, no. 4, pp. 474–486, 2014.
- [79] Z. Li, J. Gong, C. Lu, and Y. Yi, "Interactive behavior prediction for heterogeneous traffic participants in the urban road: A graph-neural-network-based multitask learning framework," *IEEE/ASME Transactions on Mechatronics*, vol. 26, no. 3, pp. 1339–1349, 2021.
- [80] Y. Lu, Y. Chen, D. Zhao, and D. Li, "Mgrl: Graph neural network based inference in a markov network with reinforcement learning for visual navigation," *Neurocomputing*, vol. 421, pp. 140–150, 2021.
- [81] E. Rehder, J. Quehl, and C. Stiller, "Driving like a human: Imitation learning for path planning using convolutional neural networks," in *International Conference on Robotics and Automation Workshops*, 2017, pp. 1–5.
- [82] Y. Peng, S.-W. Li, and Z.-Z. Hu, "A self-learning dynamic path planning method for evacuation in large public buildings based on neural networks," *Neurocomputing*, vol. 365, pp. 71–85, 2019.
- [83] S. Casas, C. Gulino, R. Liao, and R. Urtasun, "Spagnn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 9491–9497.
- [84] X. Fu, J. Zhang, Z. Meng, and I. King, "Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding," in *Proceedings of The Web Conference 2020*, 2020, pp. 2331–2341.
- [85] Z. Lu, W. Lv, Y. Cao, Z. Xie, H. Peng, and B. Du, "Lstm variants meet graph neural networks for road speed prediction," *Neurocomputing*, vol. 400, pp. 34–45, 2020.
- [86] R.-J. Wai and A. S. Prasetya, "Adaptive neural network control and optimal path planning of uav surveillance system with energy consumption prediction," *IEEE Access*, vol. 7, pp. 126 137–126 153, 2019.
- [87] X. Weng, Y. Yuan, and K. Kitani, "Ptp: Parallelized tracking and prediction with graph neural networks and diversity sampling," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4640–4647, 2021.
- [88] A. Monti, A. Bertugli, S. Calderara, and R. Cucchiara, "Dag-net: Double attentive graph neural network for trajectory forecasting," in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 2551–2558.
- [89] D. Lee, Y. Gu, J. Hoang, and M. Marchetti-Bowick, "Joint interaction and trajectory prediction for autonomous driving using graph neural networks," *arXiv preprint arXiv:1912.07882*, 2019.
- [90] J. Wiederer, A. Bouazizi, U. Kressel, and V. Belagiannis, "Traffic control gesture recognition for autonomous vehicles," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 676–10 683.
- [91] E. Tolstaya, F. Gama, J. Paulos, G. Pappas, V. Kumar, and A. Ribeiro, "Learning decentralized controllers for robot swarms with graph neural networks," in *Conference on robot learning*. PMLR, 2020, pp. 671–682.
- [92] J. Wang, W. Chi, C. Li, C. Wang, and M. Q.-H. Meng, "Neural rrt*: Learning-based optimal path planning," *IEEE Transactions on Automation Science and Engineering*, vol. 17, no. 4, pp. 1748–1758, 2020.
- [93] E. S. Low, P. Ong, and K. C. Cheah, "Solving the optimal path planning of a mobile robot using improved q-learning," *Robotics and Autonomous Systems*, vol. 115, pp. 143–161, 2019.
- [94] C. Yan, X. Xiang, and C. Wang, "Towards real-time path planning through deep reinforcement learning for a uav in dynamic environments," *Journal of Intelligent & Robotic Systems*, vol. 98, no. 2, pp. 297–309, 2020.
- [95] S. Guo, X. Zhang, Y. Zheng, and Y. Du, "An autonomous path planning model for unmanned ships based on deep reinforcement learning," *Sensors*, vol. 20, no. 2, p. 426, 2020.
- [96] H. Qie, D. Shi, T. Shen, X. Xu, Y. Li, and L. Wang, "Joint optimization of multi-uav target assignment and path planning based on multi-agent reinforcement learning," *IEEE access*, vol. 7, pp. 146 264–146 272, 2019.
- [97] C. Qu, W. Gai, M. Zhong, and J. Zhang, "A novel reinforcement learning based grey wolf optimizer algorithm for unmanned aerial vehicles (uavs) path planning," *Applied soft computing*, vol. 89, p. 106099, 2020.
- [98] Q. Yao, Z. Zheng, L. Qi, H. Yuan, X. Guo, M. Zhao, Z. Liu, and T. Yang, "Path planning method with improved artificial potential field—a reinforcement learning perspective," *IEEE Access*, vol. 8, pp. 135 513–135 523, 2020.
- [99] H. Bae, G. Kim, J. Kim, D. Qian, and S. Lee, "Multi-robot path planning method using reinforcement learning," *Applied sciences*, vol. 9, no. 15, p. 3057, 2019.
- [100] X. Lei, Z. Zhang, and P. Dong, "Dynamic path planning of unknown environment based on deep reinforcement learning," *Journal of Robotics*, vol. 2018, 2018.
- [101] A. K. Lakshmanan, R. E. Mohan, B. Ramalingam, A. V. Le, P. Veerajagadeshwar, K. Tiwari, and M. Ilyas, "Complete coverage path planning using reinforcement learning for tetromino based cleaning and maintenance robot," *Automation in Construction*, vol. 112, p. 103078, 2020.
- [102] B. Wang, Z. Liu, Q. Li, and A. Prorok, "Mobile robot path planning in dynamic environments through globally guided reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6932–6939, 2020.
- [103] M. J. Bency, A. H. Qureshi, and M. C. Yip, "Neural path planning: Fixed time, near-optimal path generation via oracle imitation," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 3965–3972.
- [104] A. H. Qureshi and M. C. Yip, "Deeply informed neural sampling for robot motion planning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 6582–6588.
- [105] Y. Li, R. Cui, Z. Li, and D. Xu, "Neural network approximation based near-optimal motion planning with kinodynamic constraints using rrt," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 11, pp. 8718–8729, 2018.
- [106] M. Saraswathi, G. B. Murali, and B. Deepak, "Optimal path planning of mobile robot using hybrid cuckoo search-bat algorithm," *Procedia computer science*, vol. 133, pp. 510–517, 2018.
- [107] B. Ichter, E. Schmerling, T.-W. E. Lee, and A. Faust, "Learned critical probabilistic roadmaps for robotic motion planning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 9535–9541.
- [108] Z. Xing and S. Tu, "A graph neural network assisted monte carlo tree search approach to traveling salesman problem," *IEEE Access*, vol. 8, pp. 108 418–108 428, 2020.
- [109] B. Fu, L. Chen, Y. Zhou, D. Zheng, Z. Wei, J. Dai, and H. Pan, "An improved a* algorithm for the industrial robot path planning with high success rate and short length," *Robotics and Autonomous Systems*, vol. 106, pp. 26–37, 2018.
- [110] A. H. Qureshi, Y. Miao, A. Simeonov, and M. C. Yip, "Motion planning networks: Bridging the gap between learning-based and classical motion planners," *IEEE Transactions on Robotics*, vol. 37, no. 1, pp. 48–66, 2020.
- [111] C. Chen, S. Hu, P. Nikdel, G. Mori, and M. Savva, "Relational graph learning for crowd navigation," in *2020 IEEE/RSJ International*

- Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10 007–10 013.
- [112] S. Aggarwal and N. Kumar, “Path planning techniques for unmanned aerial vehicles: A review, solutions, and challenges,” *Computer Communications*, vol. 149, pp. 270–299, 2020.
- [113] Á. Madridano, A. Al-Kaff, D. Martín, and A. de la Escalera, “Trajectory planning for multi-robot systems: Methods and applications,” *Expert Systems with Applications*, vol. 173, p. 114660, 2021.
- [114] E. Prianto, M. Kim, J.-H. Park, J.-H. Bae, and J.-S. Kim, “Path planning for multi-arm manipulators using deep reinforcement learning: Soft actor–critic with hindsight experience replay,” *Sensors*, vol. 20, no. 20, p. 5911, 2020.