

RRT*former: Environment-Aware Sampling-Based Motion Planning using Transformer

Mingyang Feng, Shaoyuan Li and Xiang Yin

Abstract—We investigate the sampling-based optimal path planning problem for robotics in complex and dynamic environments. Most existing sampling-based algorithms neglect environmental information or the information from previous samples. Yet, these pieces of information are highly informative, as leveraging them can provide better heuristics when sampling the next state. In this paper, we propose a novel sampling-based planning algorithm, called *RRT*former*, which integrates the standard RRT* algorithm with a Transformer network in a novel way. Specifically, the Transformer is used to extract features from the environment and leverage information from previous samples to better guide the sampling process. Our extensive experiments demonstrate that, compared to existing sampling-based approaches such as RRT*, Neural RRT*, and their variants, our algorithm achieves considerable improvements in both the optimality of the path and sampling efficiency. The code for our implementation is available on <https://github.com/fengmingyang666/RRTformer>.

I. INTRODUCTION

Motion planning in the presence of obstacles is a fundamental challenge in robotics and autonomous systems. The primary goal of this problem is to navigate a robot from a starting point to a goal point while avoiding collisions with obstacles in its environment [32], [34]. Traditional search-based approaches to this problem, such as the A* and D* algorithms [26], have been widely used due to their conceptual simplicity and efficient performance in low-dimensional spaces. However, these methods often struggle with computational complexity and scalability when applied to real-time motion planning in high-dimensional environments. As a result, there is a growing need for more advanced techniques that can handle the increased complexity and dynamic nature of modern robotic applications.

Sampling-based methods, such as Probabilistic Roadmaps (PRM), Rapidly-exploring Random Trees (RRT), and RRT* [13], [14], [16], have emerged as powerful alternatives due to their ability to handle high-dimensional spaces and complex obstacle geometries. These algorithms work by randomly sampling points in the configuration space and incrementally building a graph or tree structure that connects the start and goal points. Over the past years, sampling-based motion planning has been successfully applied in various complex engineering systems, including unmanned aerial vehicles [1], autonomous driving [31] and robot manipulators [8]. More

recently, sampling-based approaches have also been extended to handle more complex planning tasks, such as those involving high-level specifications expressed in temporal logic [4], [11], [12], [17], [19], [22], [28], [33].

A core aspect of sampling-based planning is the exploration strategy used to generate samples. Traditional methods often rely on uniform random sampling, which can be inefficient in environments with complex obstacle structures or narrow passages. To improve sample efficiency, many approaches have been developed to better guide the sampling process toward regions that are more likely to yield optimal solutions. For instance, Informed RRT [5] uses an admissible ellipsoidal heuristic to focus sampling within a subset of the configuration space that is guaranteed to contain better solutions. In [24], [27], artificial potential fields are employed to guide samples toward the goal while avoiding obstacles. However, these methods typically rely on hand-crafted heuristics, which can be challenging to adapt to new or highly complex environments.

More recently, learning-based approaches have been developed to enhance the sample efficiency in optimal path planning in complex environments. For example, in [18], a simple multi-layer perceptron (MLP) is used to estimate the cost of the path, attempting to solve the planning problem with nonlinear kinematic constraints. In [30], convolutional neural networks (CNN) is used for environmental processing to obtain a smaller sampling area. In [20], [21], graph neural networks (GNNs) are used to model the relationships between nodes to prevent collisions. In [7], the authors use point cloud information to represent the sampling area, and use pointnet to further refine the possible sampling area. In [2], using the entire map as input, a transformer is used to predict the cost of all pixels to the target point, and this is used for path planning. In [9], the transformer is used to encode environmental information and classify pixels to obtain sampling areas.

A common feature of the above methods is that they focus on environmental information and only adjust the sampling area before sampling begins, without iteratively updating it as sampling proceeds. Recently, there have been efforts to leverage information from previous sampling points to better guide future samples. For example, in [25], environmental information and previous sampling points are fused and input into a MLP to predict the next sampling point. However, this method struggles with processing complex long-distance paths because MLPs are not well-suited for handling long sequences. In [3], reinforcement learning algorithms are used to continuously iterate and interact the sampling and learning

This work was supported by the National Natural Science Foundation of China (62173226, 92367203).

M. Feng, S. Li and X. Yin are with the School of Automation and Intelligent Sensing, Shanghai Jiao Tong University, Shanghai 200240, China. (Corresponding Author: Xiang Yin) E-mail: {fmy-135214, syli, yinxiang}@sjtu.edu.cn.

processes, enabling the system to learn promising exploration directions based on the environment's structure. However, it relies on grid-based encoding of the entire workspace as input and requires the system to perceive the environment through learning rather than directly accessing environmental information. In [10], path points from the training set are encoded into discrete vectors using a transformer. The method infers the possible distribution area of each sampling point in a new environment by matching these vectors with environmental features. However, it only adjusts the sampling distribution and does not precisely determine the sampling points.

In this paper, we propose a new approach to improve sample efficiency by leveraging both environmental information and previous samples. The overall algorithm, called RRT*former, integrates the standard RRT* algorithm with the Transformer network architecture [29]. Specifically, environmental feature information is incorporated as positional encodings, enabling the algorithm to better capture the relationship between the sampling process and the environment. Sampling nodes serve as both the input and output of the Transformer, ensuring consistency between input and output. This design allows previous sampling nodes to be utilized as historical information in subsequent iterations, facilitating a better understanding of the relationships between sampling nodes. Since each sampling step is based on the current map and new information is continuously incorporated, the network can iteratively update its sampling strategy. This iterative update process is particularly crucial in dynamic environments, where conditions and information may change over time. We compare the proposed RRT*former approach with the standard RRT* algorithm and other neural network-guided sampling algorithms, such as NRRT* [30] and NIRRT* [7]. Experimental results demonstrate that our approach effectively reduces sampling time and the number of sampling nodes.

The rest of the paper is organized as follows. In Section II, we first formulate the path planning problem. Then, in Section III, we present our main algorithm, RRT*former, including its architecture and training details. In Section IV, extensive numerical experiments and simulations are provided to illustrate the effectiveness of our approach. Finally, we conclude the paper in Section V.

II. PROBLEM FORMULATION

Let \mathcal{X} denote the configuration space, where \mathcal{X}_{obs} represents the obstacle space, and $\mathcal{X}_{free} = \mathcal{X} \setminus \mathcal{X}_{obs}$ defines the free space that is navigable by the robot. Given a start state x_s , a goal state x_g , and a specific environment Env , the goal is to find a feasible path \mathcal{P} that satisfies the following conditions:

- $\mathcal{P} = \{x_i \in \mathcal{X}\}_{i=0}^n$
- $x_i \in \mathcal{X}_{free}$, for all $x_i \in \mathcal{P}$
- $Line(x_i, x_{i+1}) \in \mathcal{X}_{free}$, for all $x_i \in \mathcal{P}$
- $x_0 = x_s, x_n = x_g$

To quantify the quality of the path, we define its length using the Euclidean distance between consecutive nodes. The

total path cost is given by

$$\text{Cost}_{\mathcal{P}} = \sum_{i=1}^n \sqrt{\sum_{j=1}^d (x_i^{(j)} - x_{i-1}^{(j)})^2},$$

where n is the total number of nodes in the path and d is the spatial dimension of the configuration space \mathcal{X} . For a 2D environment, $\mathcal{X} \in \mathbb{R}^2$; for a 3D environment, $\mathcal{X} \in \mathbb{R}^3$. This cost function measures the cumulative Euclidean distance between consecutive nodes, which is typically used to evaluate the length of the path.

At each iteration step t , the configuration space is updated to reflect changes in the environment, particularly the obstacles. This update is expressed as:

$$\Delta\mathcal{X}_{obs} = f(Env(t), t, x_t),$$

where f is a sensor-based function that takes the current environmental data $Env(t)$, the current time step t , and the robot's current configuration x_t as inputs. In static environments, where obstacles do not change over time, no update occurs, i.e., $\Delta\mathcal{X}_{obs} = 0$. In dynamic environments, the function f evolves as the environment changes.

Let $\mathcal{X}_{tree}(t) = \{x_i \mid x_i \in \mathcal{X}, i < t\}$ denote the set of previously sampled nodes up to time t . Based on the principles of sample-based algorithms, \mathcal{X}_{tree} is constructed incrementally by sampling nodes from the configuration space. The process of adding new samples to \mathcal{X}_{tree} is influenced by the obstacle information at each time step. The quality of the samples generated during this process directly impacts the algorithm's convergence rate and the path's length. Therefore, it is crucial to design a sampling strategy that efficiently avoids obstacles. That is, we need to design a *Sampler* that generates the next node x_{next} based on both the current environment information and the past samples in $\mathcal{X}_{tree}(t)$. Specifically, the sampler function is expressed as:

$$x_{next} = \text{Sampler}(f, \mathcal{X}_{tree}(t)),$$

which guides the planning algorithm to explore the configuration space within the constraints of the environment.

III. RRT*FORMER ALGORITHM

In this section, we present our main sampling-based algorithm, RRT*former. Specifically, we use the standard RRT* as the basic skeleton for trajectory sampling. When sampling the next state in the RRT* algorithm, we employ a transformer architecture to generate the sampling point, leveraging both environmental information and previous sampling data. Finally, we provide a detailed description of how the proposed neural network is trained.

A. The Basic RRT* Algorithm

In this work, we adopt RRT* as the base sampling algorithm. The idea of RRT* is similar to the standard RRT algorithm. Specifically, in RRT*, the tree, denoted as \mathcal{X}_{tree} , is iteratively grown by adding nodes that not only extend toward random samples in the configuration space \mathcal{X} but

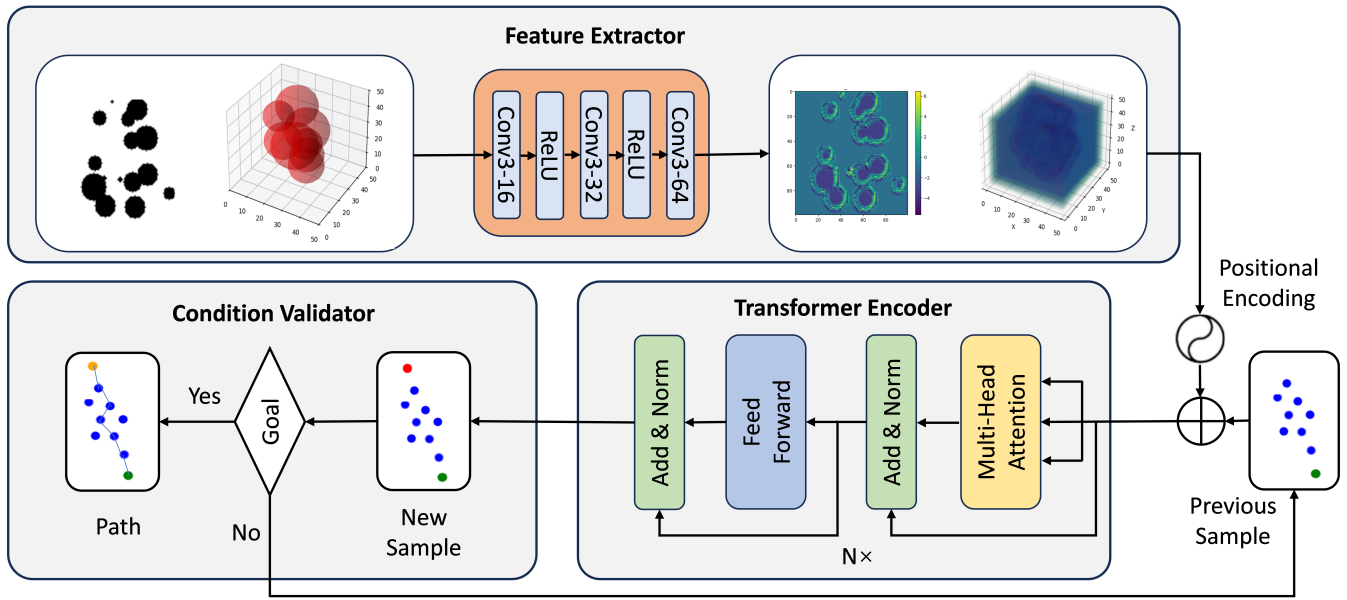


Fig. 1: Sampler Model. The Sampler consists of three parts. **Feature Extractor**: extract features from the environment using CNN. **Transformer Encoder**: generate new sample from previous sampling information and environment features. **Condition Validator**: determine when to stop sampling by checking whether the new sample is close enough to the goal.

also aim to improve the path from the starting node to the goal.

At each iteration, the following steps are executed in order:

- 1) A random node is sampled from the configuration space \mathcal{X} .
- 2) The nearest node in the tree is identified, and the tree extends toward the sampled node within a specified step size.
- 3) If the resulting path lies entirely within the free space \mathcal{X}_{free} , the new node is added to the tree.
- 4) To improve the path, RRT* performs a local optimization step where the newly added node is checked against nearby nodes in the tree. If a shorter path to the new node exists through one of these neighbors, the tree is restructured to reflect the improved path.
- 5) This process of node extension and path optimization continues until the goal node is added to the tree and the path is sufficiently optimized.

Compared to using RRT as the base algorithm, RRT* minimizes the overall path cost and ensures that the tree converges to the optimal path over time, typically resulting in a more efficient and smoother trajectory. This optimization is achieved through the rewiring of the tree as the algorithm progresses, refining the paths between nodes. The main program is outlined in Algorithm 1, where the sampling step (*Sampler*) is implemented as Algorithm 2.

B. Transformer-Based Sampler

To implement the base RRT* algorithm, the key challenge lies in efficiently sampling states from the configuration space. Here, we propose a transformer-based sampler, as illustrated in Figure 1. Specifically, the sampler model consists

Algorithm 1 Base Algorithm

Data: \mathcal{X}, x_s, x_g

Result: \mathcal{P}

```

1  $x_s \rightarrow \mathcal{X}_{tree}$  while  $x_g$  is not in  $\mathcal{X}_{tree}$  do
  /* Update knowledge */
2  $\Delta\mathcal{X}_{obs} \leftarrow f(Env(t), t, x_t)$ 
3  $\mathcal{X}_{obs} \leftarrow \mathcal{X}_{obs} \cup \Delta\mathcal{X}_{obs}$ 
  /* Take new sampled node */
4  $x_{rand} \leftarrow Sampler(f, \mathcal{X}_{tree}, \mathcal{X})$ 
5  $x_{nearest} \leftarrow Nearest(\mathcal{X}_{tree}, x_{rand})$ 
6  $x_{next} \leftarrow Steer(x_{rand}, x_{nearest}, \epsilon)$ 
  /* Check collision */
7 if CollisionFree then
8    $\mathcal{X}_{nearby} \leftarrow Nearby(\mathcal{X}_{tree}, x_{next})$ 
9    $\mathcal{X}_{tree} \leftarrow \mathcal{X}_{tree} \cup \{x_{next}\}$ 
10   $\mathcal{X}_{tree} \leftarrow Rewire(\mathcal{X}_{tree}, \mathcal{X}_{nearby}, x_{next})$ 
11 end
12 end
13  $\mathcal{P} \leftarrow GetPath(\mathcal{X}_{tree}, x_s, x_g)$ 

```

Algorithm 2 RRT* Sampler

Input: \mathcal{X}

Output: x_{rand}

```

1  $p \leftarrow$  Probability of sampling goal node
2  $r \leftarrow$  Random number between 0 and 1
3 if  $r \leq p$  then
4    $x_{rand} \leftarrow x_g$ 
5 else
6    $x_{rand} \leftarrow UniformSampling(\mathcal{X})$ 
7 end

```

of three main modules: the feature extractor, the transformer encoder, and the condition validator. The details of each component are as follows:

1) *Feature Extractor*: In the feature extractor part, we use convolutional layers to process the environment map. This module extracts useful information for sampling, such as obstacle locations and free space, from the original environmental data, which is then further processed by the transformer network.

The original environment information is represented as a 2D or 3D cost map, which is processed through CNN (for 2D maps) or 3D-CNN (for 3D maps). We employ a 3×3 convolution kernel with padding set to 1 to ensure that the spatial dimensions of the environment remain unchanged before and after convolution.

After passing through three convolutional layers, the dimensionality of the feature space is transformed into d_{model} , where d_{model} is the dimensionality of the transformer model. For example, in a 2D environment, the input to the CNN is a map of size (1, height, width), and the output is a feature map of size (d_{model} , height, width). This feature map captures essential spatial information about the environment, enabling the transformer network to make informed decisions during the sampling process.

2) *Transformer Encoder*: Once the environmental information is processed by the CNN, we further utilize a transformer model to process both the node positions and the processed environmental features. This ultimately predicts the next sampling node.

Specifically, the transformer is a network architecture composed of self-attention, multi-head attention, positional encoding, and fully connected layers. It is typically divided into two main components: the encoder and the decoder. Here, we only use the encoder part of the transformer to learn long-range dependencies between past nodes and the current environment, providing a more informed sampling node.

a) *Input Processing*. : We use the previous sampling nodes \mathcal{X}_{tree} as the original input to the transformer model. Since the number of previous sampling nodes varies each time, we first pad the sequence to a fixed maximum length. These nodes are then embedded into a higher-dimensional space using a simple linear transformation, uniformly mapping them to d_{model} dimensions. For example, in a 2D environment, we input a series of nodes with shape (2, 1) and obtain embedded nodes of shape (2, d_{model}).

b) *Positional Encoding*. : For the positional encoding of the transformer model, we take the features obtained by the Feature Extractor as input. These features are encoded using sine and cosine functions with different frequencies to effectively capture positional information. For instance, in a 2D environment, the encoding method is defined as:

$$PE(x, y, 2i) = \sin\left(\frac{x}{10000^{\frac{2i}{d_{model}}}}\right), \quad (1)$$

$$PE(x, y, 2i + 1) = \cos\left(\frac{y}{10000^{\frac{2i}{d_{model}}}}\right), \quad (2)$$

Algorithm 3 Transformer Mixed Sampler

Input: $\mathcal{X}_{tree}, Env(t)$

Output: x_{rand}

```

1  $\alpha \leftarrow$  Probability of uniform sampling
2  $r \leftarrow$  Random number between 0 and 1
3 if  $r \leq \alpha$  then
4    $x_{rand} \leftarrow UniformSampling(\mathcal{X})$ 
5 else
6   /* Feature Extraction */
7    $EnvFeature \leftarrow Conv(Env)$ 
8   /* Model Sampling */
9    $x_{rand} \leftarrow Transformer(\mathcal{X}_{tree}, EnvFeature)$ 
10 end

```

where x, y represents the position, i is the dimension index, and d_{model} denotes the dimensionality of the model.

The core of the transformer lies in the attention mechanism, whose primary computation is expressed as:

$$Attention(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V,$$

where $Q, K,$ and V denote the query, key, and value vectors, respectively. These vectors are obtained by linear transformation of the vectors composed of the environment features and the previous nodes. The term d_k represents the dimensionality of the key vectors, and the softmax function ensures that the attention weights are normalized across all keys. Through the multi-head attention mechanism and the fully connected layer, we can get the final output, i.e., new sample node with shape ($d_{space}, 1$), where d_{space} is the dimension of the sampling space, which is 2 for 2D environment and 3 for 3D environment.

3) *Condition Validator*: Finally, the condition validator serves as the third component of the *Sampler*, responsible for determining when the model should stop sampling. In transformer-based generative models, deciding when to halt the generation process is a critical consideration. For instance, in language models, this is typically managed by introducing a stop symbol or setting a maximum output length [6]. In our approach, the decision to continue sampling or stop and output the path is based on whether the newly sampled node is sufficiently close to the goal. If the sampled point is within a predefined threshold distance to the goal, the model terminates the sampling process and outputs the final path. This mechanism ensures that the algorithm efficiently converges to a solution while avoiding unnecessary computations.

With the above components detailed, we present the overall sampler as shown in Algorithm 3. Particularly, to ensure the probabilistic completeness of the sampling-based algorithm and introduce a certain degree of randomness, we incorporate a parameter α that controls the balance between uniform and transformer-based sampling. Specifically, at each sampling step, the algorithm has a probability α of performing uniform sampling and a probability $1 - \alpha$ of using transformer-based non-uniform sampling. This hybrid

approach combines the strengths of both methods: the pure randomness of uniform sampling ensures exploration of the configuration space, while the transformer-based sampling leverages environmental and historical information to guide the search more efficiently. By adjusting α , the algorithm can balance exploration and exploitation, improving both the robustness and efficiency of the planning process.

C. Training Details

To train the neural network, we construct a training set consisting of 8,000 randomly generated environments, with the optimal paths obtained using the A* algorithm. The dataset is formatted as $(\mathcal{X}|(n), x_{new}(n+1), Env(n))$, where:

- $\mathcal{X}|(n)$ represents the first n nodes on the path generated by the A* algorithm,
- $x_{new}(n+1)$ is the next node generated by the A* algorithm, and
- $Env(n)$ is the environmental information at step n .

By training the model on these optimal path planning experiences, the transformer learns a more effective sampling strategy based on historical data, which is likely to lead to faster convergence to a feasible path.

We use the Mean Squared Error (MSE) as the loss function, which calculates the error between the predicted next sampling node and the true target sampling node. The MSE is given by:

$$L_{MSE} = \frac{1}{N} \sum_{i=1}^N (x_{new}^i - \hat{x}_{new}^i)^2,$$

where x_{new} is the predicted node, \hat{x}_{new} is the actual target sampling node, and N is the batch size. During each training iteration, we use this loss function to perform backpropagation, simultaneously updating all the learnable parameters of both the Transformer and CNN.

The benefit of updating both the Transformer and CNN parameters simultaneously during backpropagation is that it enables the entire model to learn more efficiently as a unified system. This approach allows the CNN to fine-tune the extraction of environmental features while the Transformer learns to incorporate these features for sampling and path prediction. By jointly optimizing both components, the model can adapt more effectively to the relationships between the environment and node positions, leading to improved performance in tasks like path planning. Additionally, this integrated training strategy ensures that the feature extraction and decision-making processes are aligned, enhancing the model's ability to generalize to new, unseen environments.

IV. EXPERIMENTAL RESULTS

In this section, we conduct a series of experiments to evaluate the performance of the proposed algorithm. First, we provide illustrative examples in both 2D and 3D environments, demonstrating how the randomization rate α influences the algorithm's performance and highlighting its importance in balancing exploration and exploitation. Next,

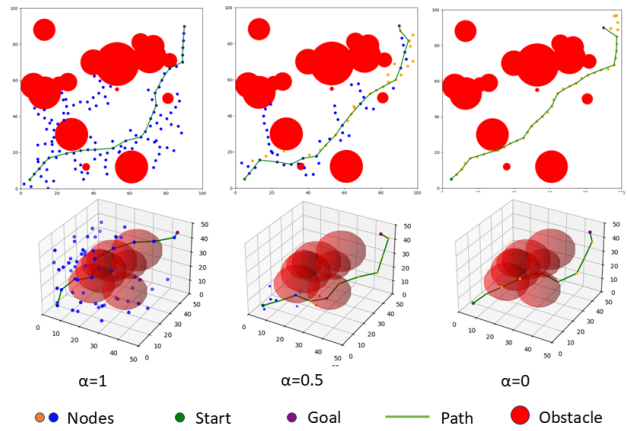


Fig. 2: Demonstration of the random tree generated under different randomization ratios α .

we compare our algorithm with existing RRT* and its variants on a large set of randomly generated environments to demonstrate its efficiency in terms of planning time and path quality. Finally, we present a case study in a simulation environment using ROS and Gazebo.

A. Experimental Settings

Our sampler model is implemented using PyTorch and trained on a single GPU (RTX 3090). The optimizer used is Adam [15] with default parameters, and the learning rate is set to 0.0001, remaining consistent throughout the training process. The transformer model is configured with the following parameters: $d_{model} = 64$, $n_{head} = 6$, and 6 encoder layers. For the 2D environment, the batch size is set to 256, and the model contains 0.52M parameters. For the 3D environment, the batch size is set to 128, and the model contains 0.56M parameters.

We consider a scenario where the algorithm must find a safe path from a starting node to an end node in a randomly generated obstacle-filled environment. Specifically,

- For the 2D environment, we generate approximately 500 scenes with a map size of 100×100 . All obstacles are circular, with the number of obstacles per scene ranging from 16 to 20 and their radii varying from 0 to 12. Using circular obstacles ensures generality, even when obstacles overlap. The step size is set to 4, the rewiring radius to 0, and the uniform sampling rate α to 0.5.
- For the 3D environment, we generate approximately 500 scenes with a map size of $50 \times 50 \times 50$. All obstacles are spherical, with the number of obstacles per scene ranging from 6 to 10. The parameter configurations, such as step size, rewiring radius, and uniform sampling rate, remain the same as in the 2D environment.

B. The Role of Hybrid Sampling

Recall that our approach employs a hybrid sampling strategy by combining transformer-based sampling and uniform sampling with a ratio α . As illustrated in Figure 2, when $\alpha = 0$, the algorithm relies purely on transformer-based

TABLE I: Statistical results under different randomization ratios α across the 500 randomly generated 2D environments.

α	Average Nodes	Success Rate(%)
0	47.02	84.14
0.5	79.70	100.00
1	275.14	100.00

sampling, and when $\alpha = 1$, it relies purely on uniform sampling. While transformer-based sampling often guides the trajectory toward the goal state more efficiently, requiring fewer sampling nodes, it may fail to achieve the planning task in some cases. Specifically, by focusing on the current environment and prior nodes, the model may overlook certain regions, potentially resulting in fewer viable paths.

To evaluate this more clearly, we conducted experiments on the randomly generated 500 scenes for the 2D environment case. We ran our RRT*former algorithm with parameters $\alpha = 0$, $\alpha = 0.5$, and $\alpha = 1$, and computed the average number of nodes in the tree and the success rate of finding a safe path from the initial state to the goal. The results are shown in Table I. The results indicate that while purely transformer-based sampling ($\alpha = 0$) generates fewer sampling nodes, it leads to a lower success rate due to its reduced randomness. On the other hand, purely uniform sampling ($\alpha = 1$) depends heavily on randomness and fails to incorporate environmental context or prior sampling information, resulting in a larger number of redundant nodes and inefficient exploration. In contrast, hybrid sampling ($\alpha = 0.5$) achieves a trade-off by leveraging existing information to reduce redundancy while preserving the randomness necessary for improving success rates.

C. Comparison with Other Methods

To further verify the efficiency of our approach, we compare the proposed RRT*former algorithm with various existing algorithms that leverage neural networks to extract environmental features. Specifically, we compare the following algorithms:

- **RRT*** (Rapidly-exploring Random Tree*)
- **NRRT*** (Neural RRT*, [30])
- **NIRRT*** (Neural Informed RRT*, [7])
- **RRT*former** (Our Algorithm)

In NRRT* and NIRRT*, we adopt the Pointnet [23] model as provided in [7].

We run each algorithm on the randomly generated 500 2D environments and 500 3D environments. For each environment, we compute the following metrics:

- **Time:** The average time required to find an initial path.
- **Nodes:** The average number of nodes explored when finding the initial path.
- **Iterations:** The average number of iterations needed to find the initial path.
- **Initial Cost:** The average length of the path found initially (before optimization).

- **Final Cost:** The average length of the path after 18 seconds of path optimization.

The statistical metrics are shown in Table II, and additional experimental details are provided in Figure 3. Based on the experiments, we draw the following observations:

Time Efficiency: One major advantage of RRT*former is its substantial reduction in computation time. Compared to traditional methods like RRT* and other learning-based algorithms such as NRRT* and NIRRT*, our method significantly shortens the time required to find an initial path. This improvement is primarily due to the model’s reduced parameter complexity and its reliance on a smaller input of sampling node data, rather than processing the entire environment representation, such as images or point clouds. This streamlined input reduces both the computational burden and the overall path-finding time, making RRT*former more efficient.

Initial Path Cost: Another significant advantage of RRT*former is the quality of the initial path it generates. The initial cost of the path found by RRT*former is consistently lower than that of other methods. This improvement stems from the transformer-based sampling strategy, which conditions its sampling process on both the current environmental state and the history of past sampling nodes.

Optimal Path and Final Cost: The final cost after optimization is another area where RRT*former shows competitive performance. Although the final cost of RRT*former in both 2D and 3D environments is slightly higher than that of RRT*, the difference is minimal. Importantly, RRT*former achieves this result much faster than both traditional RRT* and learning-based methods. It can converge to an optimal or near-optimal solution with significantly reduced computational cost.

Node Exploration Efficiency: In terms of node exploration, RRT*former demonstrates a significant advantage. It explores far fewer nodes to find a valid path compared to existing algorithms. This efficiency arises from the transformer-based sampling strategy, which learns to prioritize promising areas of the environment for sampling, rather than relying on blind random exploration.

D. Simulation Deployments

Finally, we deploy our algorithm and the model trained in a 2D random world to a TurtleBot3 Burger in the Gazebo simulator. As shown in Figure 4, the map size is $5\text{m} \times 5\text{m}$, with four static obstacles of size $1\text{m} \times 1\text{m}$ and one moving obstacle (represented as the white square) of size $0.4\text{m} \times 0.4\text{m}$. The robot’s task is to move from the initial position $(-1\text{m}, 0\text{m})$ to the final position $(-5\text{m}, 3\text{m})$ while avoiding all obstacles. The robot can detect the moving obstacle within its sensing region, shown as the purple area.

We map the $5\text{m} \times 5\text{m}$ environment into a 100×100 cost map in pixel space at a fixed ratio. After inputting the starting point, end point, and cost map into the model, we obtain a path in pixel space. This path is then reverse-mapped back to the actual simulation world and published to the robot for execution through ROS. Snapshots of the final simulation

TABLE II: Statistics on the comparison of our algorithm with other algorithms.

Scenario	Method	Time(s)	Nodes	Iterations	Initial Cost	Final Cost
2D Random Environment	RRT*	0.15	286.04	335.19	159.00	128.67
	NRRT*	0.80	202.45	314.91	151.53	132.47
	NIRRT*	1.37	164.56	211.96	147.13	132.10
	RRT*former	0.40	79.70	159.22	139.45	129.36
3D Random Environment	RRT*	0.04	86.45	98.09	105.26	77.71
	NRRT*	0.72	35.99	54.19	98.49	79.16
	NIRRT*	1.44	35.49	53.12	97.67	80.26
	RRT*former	0.29	16.65	36.81	94.82	78.27

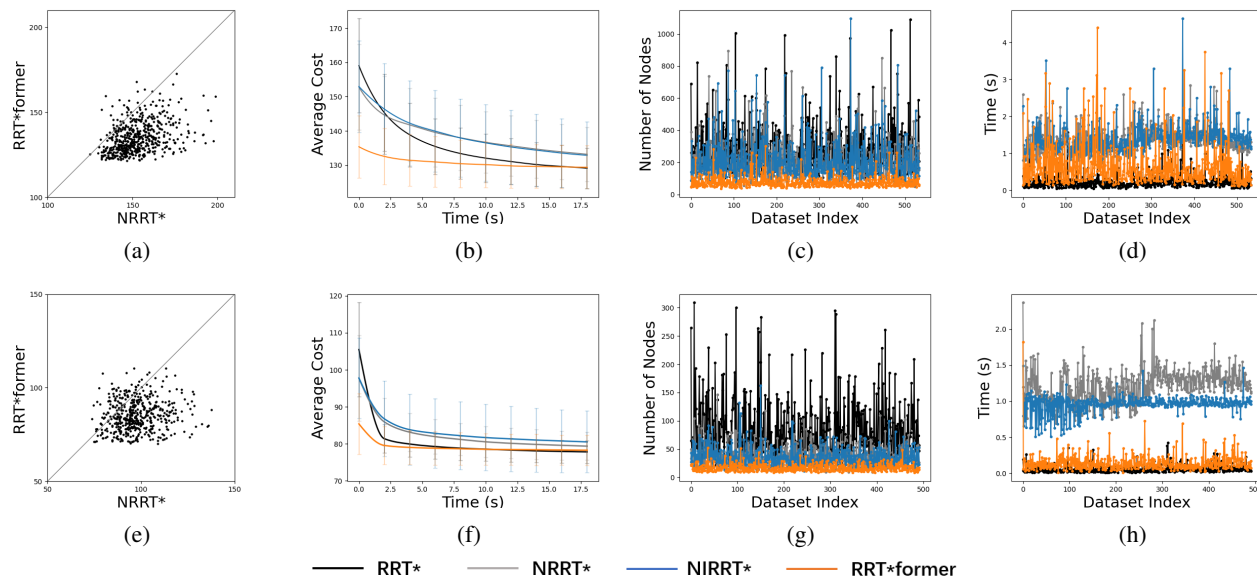


Fig. 3: Figures (a)-(d) present the results for 2D environments, while Figures (e)-(h) present the results for 3D environments. Specifically: Figures (a) and (e) compare the initial path lengths found by NRRT* and RRT*former; Figures (b) and (f) compare the path optimization speed of each algorithm after finding the initial path; Figures (c) and (g) compare the number of sampled nodes when each algorithm finds the initial path; Figures (d) and (h) compare the time taken by each algorithm to find the initial path.

results are shown in Figure 4, and the complete simulation video is available on our project website: .

As demonstrated in Figure 4, the robot successfully reaches the target state while avoiding all obstacles. Notably, since each sampling step is conditioned on both the current environmental context and the previously generated sampling nodes, the algorithm is inherently capable of avoiding dynamic obstacles in real-time. Unlike traditional random sampling methods, which treat each sampling event independently and without considering the evolving nature of the environment, this approach leverages past sampling information to generate more reasonable and forward-looking samples.

V. CONCLUSION

In this paper, we propose a new sampling-based planning algorithm, RRT*former, which integrates the standard RRT* algorithm with the Transformer architecture. Our approach

fully leverages the capabilities of the Transformer to process complex features from the environment map and learn context-dependent representations, such as the history of previously sampled states. Experimental results on randomly generated 2D and 3D environments demonstrate that, compared to existing algorithms based on RRT*, the proposed RRT*former offers considerable advantages in both path optimality and sampling efficiency. In future work, we plan to extend our approach to handle high-dimensional spaces and incorporate the kinematics of the robot to further explore the potential of this method.

REFERENCES

- [1] Cary L Butler, Reid D Smith, and Andrew G Alleyne. Sampling-based planning for guaranteed safe energy management of hybrid uav powertrain under complex, uncertain constraints. *IEEE Transactions on Control Systems Technology*, 2024.
- [2] Devendra Singh Chaplot, Deepak Pathak, and Jitendra Malik. Differentiable spatial planning using transformers. In *International conference on machine learning*, pages 1484–1495. PMLR, 2021.

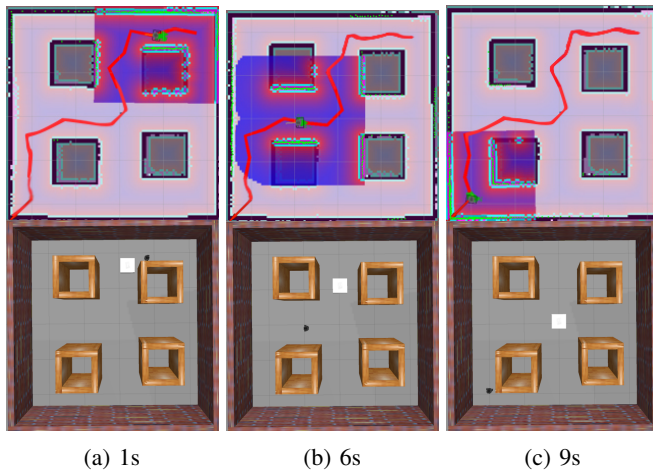


Fig. 4: Simulation Results in Gazebo: The three pictures above show the results of SLAM mapping, while the three pictures below depict the actual scenes in Gazebo. In the images, the four brown squares represent static obstacles, and the white square represents the moving obstacle. The red line indicates the path generated by our algorithm.

- [3] Binghong Chen, Bo Dai, Qinjie Lin, Guo Ye, Han Liu, and Le Song. Learning to plan in high dimensions via neural exploration-exploitation trees. *arXiv preprint arXiv:1903.00070*, 2019.
- [4] Bohan Cui, Feifei Huang, Shaoyuan Li, and Xiang Yin. Robust temporal logic task planning for multirobot systems under permanent robot failures. *IEEE Transactions on Control Systems Technology*, 2024.
- [5] Jonathan D Gammell, Siddhartha S Srinivasa, and Timothy D Barfoot. Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2997–3004. IEEE, 2014.
- [6] Lianghong Guo, Yanlin Wang, Ensheng Shi, Wanjun Zhong, Hongyu Zhang, Jiachi Chen, Ruikai Zhang, Yuchi Ma, and Zibin Zheng. When to stop? towards efficient code generation in llms with excess token prevention. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 1073–1085, 2024.
- [7] Zhe Huang, Hongyu Chen, John Pohovey, and Katherine Driggs-Campbell. Neural informed RRT*: Learning-based path planning with point cloud state representations under admissible ellipsoidal constraints. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8742–8748. IEEE, 2024.
- [8] Jinwook Huh, Boram Lee, and Daniel D Lee. Constrained sampling-based planning for grasping and manipulation. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 223–230. IEEE, 2018.
- [9] Jacob J Johnson, Uday S Kalra, Ankit Bhatia, Linjun Li, Ahmed H Qureshi, and Michael C Yip. Motion planning transformers: A motion planning framework for mobile robots. *arXiv preprint arXiv:2106.02791*, 2021.
- [10] Jacob J Johnson, Ahmed H Qureshi, and Michael C Yip. Learning sampling dictionaries for efficient and generalizable robot motion planning with transformers. *IEEE Robotics and Automation Letters*, 8(12):7946–7953, 2023.
- [11] Yiannis Kantaros and Michael M Zavlanos. Sampling-based optimal control synthesis for multirobot systems under global temporal tasks. *IEEE Transactions on Automatic Control*, 64(5):1916–1931, 2019.
- [12] Yiannis Kantaros and Michael M Zavlanos. Stylus*: A temporal logic optimal control synthesis algorithm for large-scale multi-robot systems. *The International Journal of Robotics Research*, 39(7):812–836, 2020.
- [14] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. [13] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011.
- Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Steven M LaValle. *Planning algorithms*. Cambridge university press, 2006.
- [17] Shuaiyi Li, Mengjie Wei, Shaoyuan Li, and Xiang Yin. Temporal logic task planning for autonomous systems with active acquisition of information. *IEEE Transactions on Intelligent Vehicles*, 9(1):1436–1449, 2023.
- [18] Yang Li, Rongxin Cui, Zhijun Li, and Demin Xu. Neural network approximation based near-optimal motion planning with kinodynamic constraints using rrt. *IEEE Transactions on Industrial Electronics*, 65(11):8718–8729, 2018.
- [19] Ruijia Liu, Ancheng Hou, Xiao Yu, and Xiang Yin. Zero-shot trajectory planning for signal temporal logic tasks. *arXiv preprint arXiv:2501.13457*, 2025.
- [20] Ruijia Liu, Shaoyuan Li, and Xiang Yin. NNgtL: Neural network guided optimal temporal logic task planning for mobile robots. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10496–10502. IEEE, 2024.
- [21] Wansong Liu, Kareem Eltoumy, Sibio Tian, Xiao Liang, and Minghui Zheng. Kg-planner: Knowledge-informed graph neural planning for collaborative manipulators. *IEEE Transactions on Automation Science and Engineering*, 2024.
- [22] Xusheng Luo, Yiannis Kantaros, and Michael M Zavlanos. An abstraction-free method for multirobot temporal logic optimal control synthesis. *IEEE Transactions on Robotics*, 37(5):1487–1507, 2021.
- [23] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [24] Ahmed Hussain Qureshi and Yasar Ayaz. Potential functions based sampling heuristic for optimal path planning. *Autonomous Robots*, 40:1079–1093, 2016.
- [25] Ahmed Hussain Qureshi, Yinglong Miao, Anthony Simeonov, and Michael C Yip. Motion planning networks: Bridging the gap between learning-based and classical motion planners. *IEEE Transactions on Robotics*, 37(1):48–66, 2020.
- [26] Anthony Stentz. The focussed D* algorithm for real-time replanning. In *IJCAI*, volume 95, pages 1652–1659, 1995.
- [27] Zaid Tahir, Ahmed H Qureshi, Yasar Ayaz, and Raheel Nawaz. Potentially guided bidirectionalized rrt* for fast optimal path planning in cluttered environments. *Robotics and Autonomous Systems*, 108:13–27, 2018.
- [28] Cristian Ioan Vasile, Xiao Li, and Calin Belta. Reactive sampling-based path planning with temporal logic specifications. *The International Journal of Robotics Research*, 39(8):1002–1028, 2020.
- [29] 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.
- [30] Jiankun Wang, Wenzheng Chi, Chenming Li, Chaoqun Wang, and Max Q-H Meng. NeuralRRT*: Learning-based optimal path planning. *IEEE Transactions on Automation Science and Engineering*, 17(4):1748–1758, 2020.
- [31] Jiankun Wang, Tingguang Li, Baopu Li, and Max Q-H Meng. GMR-RRT*: Sampling-based path planning using gaussian mixture regression. *IEEE Transactions on Intelligent Vehicles*, 7(3):690–700, 2022.
- [32] Xiang Yin, Bingzhao Gao, and Xiao Yu. Formal synthesis of controllers for safety-critical autonomous systems: Developments and challenges. *Annual Reviews in Control*, 57:100940, 2024.
- [33] Xinyi Yu, Xiang Yin, Shaoyuan Li, and Zhaojian Li. Security-preserving multi-agent coordination for complex temporal logic tasks. *Control Engineering Practice*, 123:105130, 2022.
- [34] Jianing Zhao, Keyi Zhu, Mingyang Feng, Shaoyuan Li, and Xiang Yin. No-regret path planning for temporal logic tasks in partially-known environments. *The International Journal of Robotics Research*, 2025.