Autonomous Robot Motion Planning Technology Based on Diffusion Model

Hanbin Sun

{2209008209@wya.top }

Beijing World Youth Academy, Beijing, 10010, China

Abstract. In recent years, the diffusion model has emerged as a powerful type of generative AI and has demonstrated remarkable potential in motion planning tasks. This study presents a systematic review on the application of diffusion models in the motion planning of autonomous robots. The review covered the application of diffusion models across five key domains: Static environments, Dynamic environments, Single-robot systems, Multiple-robot systems, and Human-robot interaction scenarios. For each of these categories, the paper presents key methodologies used, analyzes how the application of diffusion systems address core challenges, and discusses their advantages and limitations. Through surveying existing papers, the review highlights several challenges faced by current applications of diffusion models, such as reliance on high quality data and low computational efficiency. Moreover, the review discusses possible solutions to these challenges, highlighting potential research directions in the application of diffusion models.

Keywords: Motion Planning, Path Planning, Autonomous Robots, Diffusion Model

1 Introduction

With the advent of the automation era, the application scenarios of automatic robot systems are becoming increasingly extensive, ranging from precise industrial manipulator to intelligent service robots. In which, motion planning has consistently served as the key for ensuring the safe and efficient operation of autonomous systems. Over the past decades, the research on motion planning, evolving from the traditional search-based and sampling-based methods to the emerging learning-based paradigm, has continually pushed the boundaries of robotic autonomy. The application of autonomous robots have thus become increasingly widespread, gradually entering the public eye.

Between 1970 and 1990, autonomous robots were mainly applied in simple low-dimensional environments under standardized conditions, with examples including industrial manipulators and autonomous agents for structured indoor environments. In which, a trend toward AI-based motion planning is obvious, and the Diffusion Model appears to be one of the greatest focuses on this trend. In these early applications, motion planning is mainly achieved through discrete-based approaches. These approaches, exemplified by Dijkstra, A*, and D*, assumed the

ICIAAI 2026, May 22-24, Singapore, Singapore

DOI 10.4108/eai.22-5-2026.2365337

external environment as a graph of nodes or a grid of cells. Each of the nodes are assigned with a movement cost according to metrics such as distance, and a path with the minimum cost is calculated [1]. Between 1990 and 2000, sampling-based approaches of motion planning, exemplified by the probabilistic roadmap method (PRM) and the rapidly-exploring random tree (RRT), begin to emerge [2]. These approaches enabled robot agents to navigate in not only structured obstacle environments but also in irregular high-dimensional environments, such as outdoor navigation in irregular landscapes. In sampling-based approach, points are selected randomly in the free spaces of the environment as well as on the boundary of obstacles. Points that are navigable to each other are then connected to form a graph, where the optimal path is finally selected [1]. Around the 2010s, approaches based on machine learning illuminates a new path toward motion planning. This approach achieves motion planning by gradually developing or learning the best solution instead of defining the logic of robot motion through math and algorithms. Through development and training with enough data, a path could theoretically be generated at any moment for any environment. The integration of Machine learning in motion planning, therefore, has enabled autonomous robots to navigate in unseen unstructured environments.

In the present, researches on motion planing are aiming at integrating autonomous robot systems into human society and into people's everyday life. According to the International Federation of Robotics, the emergence of generative AI leads the transition from rule-based automation to intelligent, self-developing systems. This enabled robots to deal with unplanned situations and efficiently generate high-quality paths in real-world scenarios. The newly emerged Diffusion model, as a class of generative AI, has recently demonstrated remarkable potential in motion planning tasks. Unlike traditional discriminative models that learn decision boundaries, diffusion models learn the underlying distribution of feasible trajectories through an iterative denoising process. In which, they are trained to gradually recover clean trajectory data from a random noise by reversing a fixed forward diffusion process [3]. In recent years, a growing body of work has emerged applying diffusion models to various kinds of motion planning, ranging from static environment navigation to dynamic crowd interactions, from single-robot systems to multi-robot coordination, and from pure geometric planning to language-instructed behavior generation. This surge of interest motivates a systematic review of how diffusion models are reshaping the landscape of robot motion planning.

In this paper, the recent advances in motion planning of autonomous robots on diffusion models is comprehensively reviewed. Section 2 presents some common approaches on motion planning, and section 3 introduces the fundamental principles of diffusion models. Then, section 4 surveys existing papers across five key application domains: Systems in static environment, Systems in dynamic environments, Single-robot systems, Multiple-robot systems, and Human-robot interaction scenarios. For each of these categories, their advantages and limitations is discussed. Finally, section 5 discusses the current development of diffusion model in motion planing, identify open challenges, and outline promising directions for future research.

2 Common approach on motion planning

Motion planning refers to the process of finding a collision free path used for the robot to complete a specific task while conforming with the constraints of kinematics, dynamics and energy consumption. Throughout the history of motion planning, common approaches can be classified into five different categories: Discrete-based Planning, Gradient-based Planning, Sampling-based Planning, Bio-inspired Planning, and Planning with Machine Learning. In which, the first three can be grouped as Traditional algorithmic Approach, while the rest can be grouped as AI-based approach.

2.1 Algorithmic approaches

In traditional algorithmic approaches, movement of robots are pre-coded and mathematically defined. There are three algorithmic approaches: Discrete-based Planning, Gradient-based Planning, Sampling-based Planning.

Discrete-based Planning divides an environment into a grid of cells or graph of nodes. Each node is assigned with a movement cost according to a specific function, and a path with the minimum cost is computed as the optimal path. The first well-known path planning algorithm, The Dijkstra's algorithm, utilizes this approach, together with its two improved versions A* and D*. Although gaining initial success, discrete based planning face challenges when handling large environments due to high computational costs [1]. Also, it isn't always feasible to directly generate explicit discrete graphs in complex high-dimensional environments.

Sampling-based Planning arouse to solve the issue with discrete based planning in high-dimensional complex environments. Instead of modeling the environment, it selects points from the environment and connects points that are navigable between each other to form a network of paths. The optimal path is then select within the network. PRM and RRT are all well-known examples of Sampling based planning. Sampling based method like RRT and PRM are suitable for complex, high dimensional environments, as they do not require an explicit model of the environment [4]. They are also generally computationally efficient for the same reason. However, the path quality generated for sampling-based planning are highly dependent on the sampling strategy used, in which inappropriate sampling method may result in suboptimal paths.

Gradient-based Planning emerges in about the same time as sampling-based planning. It transfers the environment into a virtual landscape, where low-altitude valleys lead to the target and obstacles are represented by high-altitude areas. Robot is directed to follow the path with the steepest descend, thereby reaching the target. The most prominent example of this method is the Artificial Potential Field (APF). The special obstacle avoidance method of Gradient based planning allowed it to handle dynamic obstacles effectively while always ensures collision avoidance. However, gradient-based planning often faces two limitations: susceptibility to local minima, and poor performance in traversing narrow passages and handling multi-goal scenarios [1].

2.2 AI-based approaches

Algorithmic approaches provide robot systems with reliable pathfinding capability but lack the robustness and adaptability to deal with complex and unplanned scenarios. AI-based approach tackles this challenge. In AI-based approaches, motion planning is achieved through adaptive and learning-based approaches, instead of directly defining the logic of robot motion through math and algorithms.

Bio-inspired algorithms are the earliest AI-based approach widely used in automatic robot systems. They draw inspirations from the processes, phenomenas, and behaviors of the natural world and perform simulation to solve complex problems. The bio-inspired algorithms include the genetic algorithm and the particle swarm optimization. The genetic algorithm simulate the process of natural selection to develop the most promising solution. The particle swarm optimization mimics the phenomenon of swarm intelligence to navigate toward the target. As solutions are developed instead of algorithmically defined, bio-inspired algorithms are more adaptable to unplanned environments. However, their performance is sensitive to parameter settings, and there're potential for slower convergence in large-scale environments.

Machine learning are the most popular AI-based approach, applied in many state-of-the-art motion planning models. Machine learning models that are included in motion planning encompasses multiple types, such as Supervised learning, imitation Learning, and Reinforcement learning, but basically can be categorized into two main categories: discriminative models and generative models. Discriminative models, such as traditional reinforcement learning models, learn decision boundaries that directly map inputs to outputs, enabling them to efficiently predict and generate feasible paths. However, they struggle with multimodal problems where multiple equally valid solutions exist, often collapsing to an average solution that satisfies none of the individual modes [5]. This limitation has motivated the exploration of generative approaches, which learns the underlying distribution of all possible solutions. Among which, the diffusion model has emerged as a particularly promising paradigm for robot motion planning.

3 The Diffusion Model

The Diffusion model, with full name Denoising Diffusion Probabilistic Model (DDPM), which is an innovative class of generative AI. It was originally designed and created for high-quality image generation, but had demonstrated outstanding capabilities when applied in other fields such as motion planning. The core idea of DDPM is to convert random gaussian noise into clear data pattern by training a neural network to perform an iterative denoising process. As shown in Fig. 1, the operation of the diffusion model involves two separate process: the forward process and the backward process [3].

The forward process, also called the destruction process, draws inspiration from the physical random diffusion process that particles undergo when transforming from a low-entropy state to a high-entropy state. Starting from a clean data sample X_0, the model takes multiple Markovian steps to destroy the structure of X_0. During each step, a small amount of gaussian noice is added to the data. After T steps, the data will be turned into pure random noise X_T. Due to the additive nature of gaussian distributions, X_T can be obtained directly from the initial data X_0 without iterating through all T steps individually.

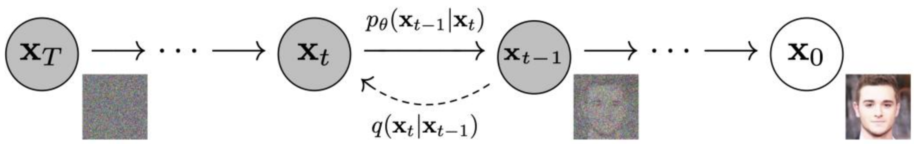

Fig. 1. The visualization of the DDPM model [3].

In the backward process, the model learns to reverse the destruction process. A neural network is trained to predict and remove the noise that was added in the forward process. During training, the network learns through a cost function that compares the predicted noise with the actual noise added. During the actual generation process, starting from random noise, the model iteratively denoises the sample over T steps, gradually recovering a structured output that resembles the initial data distribution X_0. This iterative refinement enables DDPM to generate high-quality and diverse samples.

However, the standard DDPM requires hundreds of denoising steps to generate a single sample, which can be infeasible for real-time robotics applications where frequent changes occur in the environment. To address this limitation, Denoising Diffusion Implicit Models (DDIM) were proposed as a technique to enhance efficiency and reduce computational stress [6]. DDIM introduces a non-Markovian sampling procedure that allows the generation process to skip numerous amounts of intermediate steps while mostly maintaining generation quality. This advancement made diffusion models practical for real-time robot motion planning, where high reaction speed and low latency is required.

4 Application of Diffusion Models in Autonomous Robot Systems

In this section, a comprehensive survey of the existing applications of diffusion models in motion planning will be presented. This paper surveyed across five key application domains: Systems in static environment, Systems in dynamic environments, Single-robot systems, Multiple-robot systems, and Human-robot interaction scenarios.

4.1 Systems in static environment

Algorithmic approaches provide robot systems with reliable pathfinding capability but lack the robustness and adaptability to deal with complex, unplanned scenarios. AI-based approach tackles this challenge. In AI-based approaches, motion planning is achieved through adaptive and learning-based approaches, instead of directly defining the logic of robot motion through math and algorithms.

Autonomous robots have long been applied to a wide range of industrial processing tasks, including polishing, robotic painting, and spray coating, where robots must follow reasonable trajectories across irregular component surfaces. To enhance the efficiency of performing such tasks, Chen et al. proposed 3D-CovDiffusion [7], a diffusion-based framework that generates long, smooth trajectories adaptable to various geometrical shapes without the need of manual redesign or extensive parameter tuning. By learning the distribution of feasible paths directly from data, 3D-CovDiffusion enables high surface coverage across different industrial applications, as shown in Fig. 2(a). Nevertheless, the approach faces limitations due to the scarcity of training data, which can lead to low precision when handling complex surfaces such as deep cavities and narrow rims.

The application of diffusion models in static environment also extend into indoor robot services, where robots must efficiently perform dedicated tasks in dense obstacle environments. Arthur et al. designed a system that combined diffusion-based warm-starting with a collision-aware model predictive controller [8]. This method perfectly offsets the limitations of both diffusion and MPC through first generating a sample path with diffusion and then use MPC for optimization, which allows it to achieve quicker and more robust path generation for robot arms. Current, the research is conducted on a 2D-based environment, as shown in Fig. 2(b), but it is promising for extending into 3D in the future.

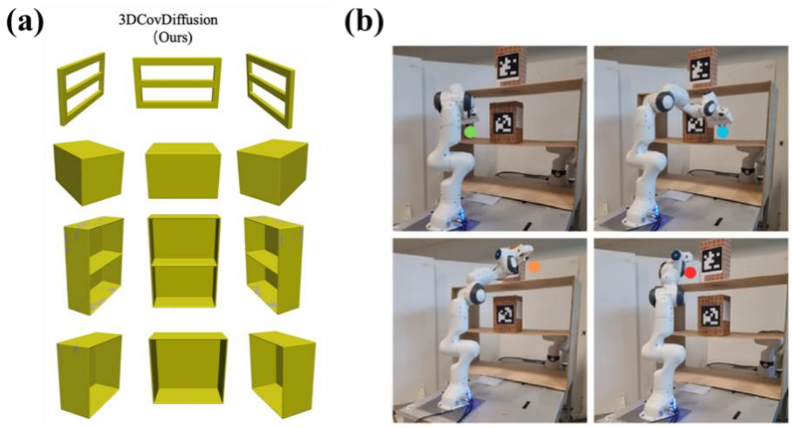

Fig. 2. (a) Simulation of painting coverage of 3D CovDiffusion when painting various furnitures [7]. (b) A robot arm, utilizing Collision-free MPC with diffusion-based warm start, effectively reaching end-effector positions while avoiding obstacles [8].

4.2 Systems in dynamic environment

Autonomous vehicles have been one of the commonly discussed topics within the last few years. For real-world Autonomous vehicles, it is a challenge to ensure both stringent collision avoidance and real-time planning efficiency. Mao et al. proposes a rapid and safe trajectory planning (RSTP) framework [9] that utilize a diffusion planner to transform raw sensor inputs directly to executable trajectories, achieving fast inference without post-hoc optimization. In which, a diffusion composition mechanism that combines separately trained models was introduced to generalize to unseen hybrid scenarios, such as static and dynamic obstacles combined, together with a lightweight rule-based safety filter that selects the best trajectory to ensuring safe deployment even when individual diffusion samples malfunctions. When tested on a real F1TENTH vehicle, as shown in Fig. 3, this framework ensured safe, real-time trajectory planning with a superior 0.21s mean planning time [9]. Nevertheless, the performance of diffusion composition depends heavily on tuning the weights assigned to each mode, in which wrong weights can increase failure rates significantly.

Fig. 3. A F1TENTH vehicle navigating through dynamic environments using the RSTP method. Green denotes vehicle trajectory, while red denotes the trajectory of dynamic obstacle [9].

Furthermore, as automatic robot system integrates into human society, safely navigating crowds become one of the basic requirements. Samavi et al. introduces SICNav-Diffusion [10], a system that achieves bidirectional coupling between robot planning and human prediction in crowded dynamic environments. The team proposes the Joint Motion Indeterminacy Diffusion (JMID) prediction model, a diffusion model utilized to predicts the joint future trajectories of all humans in a scene simultaneously. Predictions are then passed into a Bilevel Model Predictive Control (MPC) framework, where the upper level plans the robot's path and the lower level refine the predictions to ensure non-collision. This combination enables robot agents to simultaneously refine human trajectory predictions while planning their own path, resulting in safer, more efficient navigation through human crowds without freezing or colliding.

4.3 Single-robot systems

Single robot scenarios are very common in the application of autonomous robots. In which, the application of diffusion model enhanced the efficiency and feasibility of motion planning. Cai proposes the Navigation Diffusion Policy (NavDP) [11], an end-to-end transformer-based diffusion policy with dual-head architecture that jointly performs trajectory generation via diffusion and trajectory evaluation via a critic network. This system achieves zero-shot sim-to-real transfer across different robot types and environments, as shown in Fig. 4, without the need for any real-world fine-tuning, allowing efficient trajectory planning across multiple single-robot scenarios. Nevertheless, the zero-shot generalization depends heavily on the diversity of training data. The overall performance will degrade significantly if real-world robot height or camera parameters fall outside the randomized range used during training.

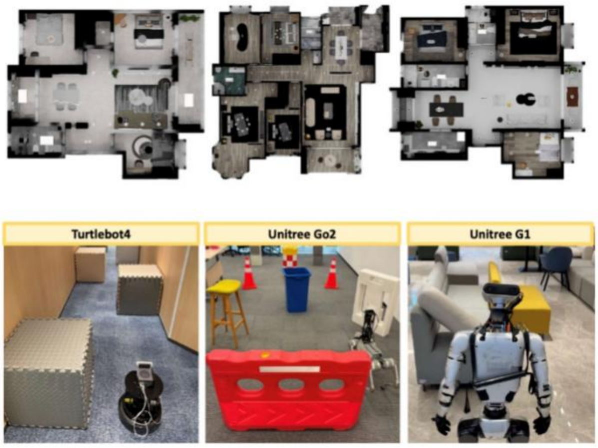

Fig. 4. Application of NavDP on various robot types [11].

Furthermore, Balim proposes the Model Predictive Diffuser (MPD) [12], a compositional diffusion framework that combines a diffusion planner, a diffusion dynamics model and a ranker for selection. In which, the planner generates task-aligned trajectories while the dynamics model ensure that the trajectory generated are dynamically feasible. The framework introduces an alternating sampling scheme where the planner and dynamics model update the trajectory in turns during denoising, progressively achieving high dynamic feasibility while ensuring task completion. Nevertheless, the framework still require the planner and dynamics model to be trained on compatible data distributions.

4.4 Multi-robot systems

In cooperative tasks involving multiple robot agents, individual agents need to navigate efficiently while preventing collision with other agents. Liang proposed a framework that integrates constrained optimization with diffusion models for multi-agent path finding in continuous spaces [13]. This enabled the simultaneous generation of trajectories for all agents while considering collision avoidance and kinematic constraints, as shown in Fig. 5. The system embedded constraints directly into the sampling process of diffusion to ensure the direct generation of feasible paths. An augmented lagrangian method is also introduced to the framework to efficiently handle computationally expensive scenarios.

In addition, in multi-robot scenarios, diffusion models are also combined with other earlier obstacle-avoiding mechanisms to construct robust motion-planning systems. Chae proposed Language-Conditioned Heat-Inspired Diffusion (LCHD) [14], an end-to-end vision-based framework that generates collision-free trajectories for multiple robots directly from raw RGB images and natural language instructions. The system combines diffusion with a heat-inspired gradient-based approach, where the target is viewed as a heat source and obstacles are viewed as insulators. By integrating heat-inspired diffusion with semantic priors based on Contrastive Language–Image Pre-training (CLIP), the system achieves the accurate execution of language commands while responding flexibly to the physical environment: when a specified goal is blocked by obstacles, the robot automatically redirects to an accessible alternative that matches the semantic intent.

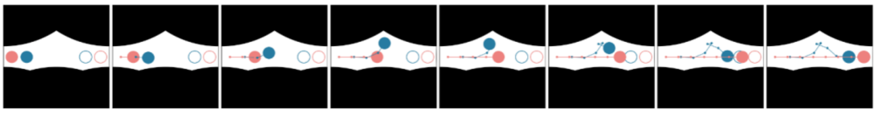

(a) Narrow Corridor 1.

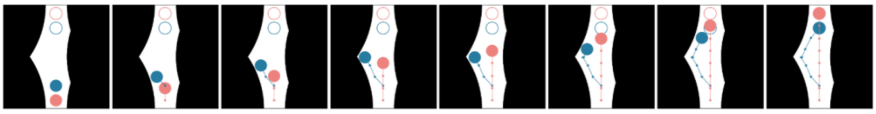

(b) Narrow Corridor 2.

Fig. 5. Collision-free trajectories generated by constrained optimization with diffusion models in narrow corridors. Solid circle represents robot agents [13].

4.5 Human-robot interaction scenarios

Human-robot interaction gained significant importance as autonomous robot systems enters human society. Hu proposed ComposableNav, a diffusion-based motion planner that handles several instruction specifications simultaneously, such as overtake pedestrian and stay on right side [15]. It achieves this by summing the predicted noise from multiple diffusion models, each handling one specification, during the denoising process, generating trajectories that simultaneously satisfy multiple instruction specifications. This system allows autonomous robot systems to more efficiently function within human society, where language communication serves as the main human-robot interaction. An example demonstration of the planner in real-world is shown in Fig. 6.

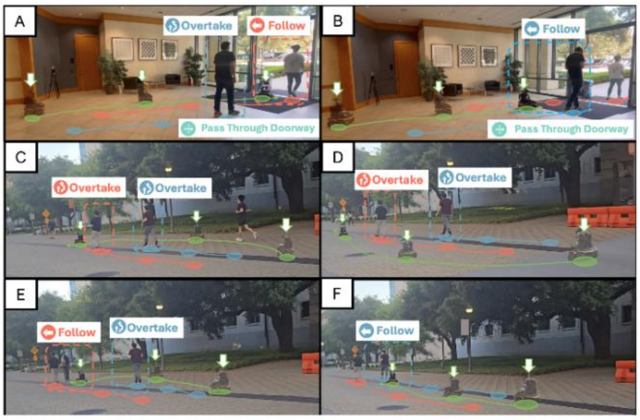

Fig. 6. Robot agents using ComposableNav navigating in real world human-interaction scenarios, handling several instruction specifications simultaneously [15].

In addition to obeying instructions, Human-robot interaction also involves safely navigating through human crowd whist completing tasks. Responding to this challenge, the SICNav-Diffusion enables robots to simultaneously refine human trajectory predictions while planning their own path, achieving safe and efficient navigation through dense human crowd without frequently freezing or colliding [10].

5 Limitations and future prospects

Diffusion model has demonstrated immense potential in the field of robotics, enabling breakthroughs in multiple areas ranging from coverage path planning in static scenarios to human-robot interaction in dynamic scenarios. By reviewing current applications of diffusion model in motion planning, some common challenges can be identified, and these challenges provide valuable insights for future research directions.

5.1 Reliance on high quality data

Although diffusion models offer exceptional robustness in motion planning, their performance relies heavily on the scale and quality of the training data they received. However, collecting high-quality data, including expert demonstrations and diverse real-world scenarios, is both costly and time consuming. Scarcity of high-quality training data may lead to performance degradation in complex environments. For 3D-CovDiffusion, for instance, the scarcity of training data had led to registration errors when the model encounters complex geometric surfaces.

To deal with this challenge, future research can explore more efficient methods of data utilization. To mitigate the issue of data scarcity and reduce the demand for specific data for downstream tasks, researches can explore ways of utilizing the basic VLM model pre-trained on internet level data to provide comprehensive prior knowledge for diffusion planners. In addition, the method utilized by NavDP, which is to generate numerous amount of data through a highly efficient generation pipeline, can also be used for reference, but to a further extent. To be specific, researches could focus on creating a more realistic and diversified simulation data generation pipeline, and explore effective simulation to reality adaptation technology to remove the gap between simulation and reality.

5.2 Sensitive to data distribution

As an AI-model, diffusion models are highly sensitive to the distribution of training data, which limits their generalizability. As shown by NavDP, once the parameters of the real robot, such as camera height, exceed the trained distribution range, the success rate will drop sharply. MPDiffuser, as another example, uses a complex alternating sampling mechanism between a diffusion planner and a diffusion dynamic model to ensure the feasibility of path samples and enhance generalizability. However, precise coordination of training data is required between the two models, and significant degradation in performance will occur if training data is mismatched. The RSTP model, on the other hand, tried to ensure feasibility and generalizability by combining multiple separately trained models. However, when dealing with unseen combined scenarios, the model’s performance depends heavily on the fine tune of weights assigned to the individual models, and incorrect weights may lead to failure in generalization.

Despite current shortcomings, composition-based diffusion still have the potential to be a promising way to enhance the generalizability of diffusion models. The composition methods of current applications such as the RSTP and the ComposableNav are still relatively elementary, being sensitive to weights and face challenges when handling complex combinations. Future studies can therefore explore more robust and principled compositional methods, such as energy-based models or factor graphs, to achieve more stable trajectory generation.

5.3 Low computational efficiency

Computational efficiency is crucial for autonomous robots in dynamic real-world environments, as promptness of reaction determines the safety of both human individuals and the robots themselves. Diffusion models, however, often require hundreds of steps of denoising to generate high-quality samples, which poses challenges for planners on achieving the real-time performance required for robot control in dynamic real-world environment. The SICNav-Diffusion, for instance, compromised the quality of predictions to use the accelerated DDIM instead of the standard DDPM, in order to reach the desired control frequency of 10 Hz. The NavDP paper had also acknowledged that, despite reducing the number of denoising steps to 10, there is still delays when compared to the single-step prediction method [11].

To enhance computational efficiency and ensure the timeliness of autonomous robots, future research can focus on developing more advanced fast samplers, transcending DDIM, that maintains sample quality while reducing the number of denoising steps to single digits. Furthermore, a possible way to accelerate the overall motion-planning process and reduce computational stress is to combine a diffusion base model with a traditional global planner or learning-based potential planner. In which, the diffusion model will focus only on short-term local details, while the long-term dependencies and the understanding of global tasks will be handled by the other planner. This hierarchal architecture could possibly reduce problem complexity and enhance the overall real-time performance of autonomous robots.

6 Conclusion

In conclusion, this study presents a comprehensive review of recent advances in motion planning of autonomous robots on diffusion models. The study first summarizes the history of motion planning in autonomous robots, from the basic discrete-based algorithms to the advanced AI-based planners, demonstrating that the integration of autonomous robot systems into human society is the main trend. Next, the study presents an explanation of the basic operating principles of DDPM and DDIM, laying the knowledge foundation for the afterward review. Then, the study analyzes state-of-the-art applications of diffusion models in motion planning in static environments, dynamic environments, single-robot systems, multiple-robot systems, and human-robot interaction scenarios, demonstrating that the application of diffusion models has significantly improved the overall robustness and generality of motion planning. Finally, the study points out that existing diffusion-based planners still face challenges in generalizability and computational efficiency, due to their reliance on high quality data, sensitiveness to data distribution, and long generation procedure. In the future, to further push the boundaries of motion planning with diffusion models, studies can be conducted to explore more efficient methods of data utilization and to develop more advanced lightweight architecture designs, in order to enhance generalizability and accelerate the motion-planning process. The findings of this review may serve as a solid reference for both academic scholars and robot practitioners. It offers a clear presentation of existing capabilities of diffusion models, and points out future opportunities to harness generative AI for more adaptive and intelligent autonomous systems.

References

[1] AbuJabal, N.: A Comprehensive Study of Recent Path-Planning Techniques in Dynamic Environments for Autonomous Robots. Sensors, pp: 8089–8089 (2024)

[2] Orthey, A.: Sampling-Based Motion Planning: A Comparative Review. Annual Review of Control, Robotics, and Autonomous Systems, vol. 7, no. 1, Annual Reviews (2023)

[3] Ho, J.: Denoising Diffusion Probabilistic Models. arXiv (2020)

[4] Waga, A.: A survey on autonomous navigation for mobile robots: From traditional techniques to deep learning and large language models. J. King Saud Univ. Comput. Inf. Sci. pp. 198 (2025)

[5] Kasa, S. R.: Generative or Discriminative Revisiting Text Classification in the Era of Transformers. arXiv (2025)

[6] Song, J.: Denoising Diffusion Implicit Models. arXiv (2022)

[7] Chen, C.: 3D-CovDiffusion: 3D-Aware Diffusion Policy for Coverage Path Planning. arXiv (2025)

[8] Haffemayer, A.: Warm-Starting Collision-Free Model Predictive Control with Object-Centric Diffusion. arXiv (2026)

[9] Mao, W.: Rapid and Safe Trajectory Planning over Diverse Scenes through Diffusion Composition. arXiv (2025)

[10] Samavi, S.: SICNav-Diffusion: Safe and Interactive Crowd Navigation with Diffusion Trajectory Predictions. IEEE Robotics and Automation Letters, Institute of Electrical and Electronics Engineers, pp. 1–8 (2025)

[11] Cai, W.: NavDP: Learning Sim-To-Real Navigation Diffusion Policy with Privileged Information Guidance. arXiv (2025)

[12] Balim, H.: Model-Based Diffusion Sampling for Predictive Control in Offline Decision Making. arXiv (2025)

[13] Liang, J.: Multi-Agent Path Finding in Continuous Spaces with Projected Diffusion Models. arXiv (2025)

[14] Chae, J.: Multi-Robot Motion Planning from Vision and Language Using Heat-Inspired Diffusion. arXiv (2025)

[15] Hu, Z.: ComposableNav: Instruction-Following Navigation in Dynamic Environments via Composable Diffusion. arXiv (2025)