



Intelligent Path Planning Algorithms for Autonomous Mobile Robots in Dynamic Environments

Taneesha Fernandes-Pereira

Department of Civil and Environmental Engineering, Basra Institute of Business Technology, Iraq
taneesha.fernandes.pereira@bibt-iq.org

Peer Review Information

Submission: 03 Jan 2020

Revision: 18 Feb 2020

Acceptance: 06 March 2020

Keywords

Autonomous mobile robots,
 Path planning, Dynamic
 obstacle avoidance,
 Dynamic Window Approach,
 Reinforcement learning

Abstract

Autonomous mobile robots operating in warehouses, hospitals, campuses, and service environments must plan collision-free motion while the surrounding geometry changes because of pedestrians, vehicles, and other robots. This paper presents a methodology for evaluating intelligent path planning algorithms in dynamic environments and develops a hybrid global-local planning framework that combines graph-based route generation, short-horizon obstacle prediction, and reactive velocity selection. The study considers A* replanning, artificial potential fields, a Dynamic Window Approach (DWA)-inspired local planner, and a predictive hybrid planner. A 30×30 occupancy-grid benchmark containing fixed structures and moving obstacles is used to test 50 randomized scenarios under common start-goal conditions. Performance is assessed using success rate, path length, turning events, clearance, and computation time. The results show that the predictive hybrid planner achieves the highest completion rate (92%) while preserving a path length close to repeated A* replanning. Pure potential-field and local DWA-inspired strategies are computationally light but show more failures and longer paths because of local minima or limited global guidance. The methodology demonstrates that intelligent navigation in dynamic environments benefits from combining global optimality information with local motion feasibility and short-term prediction rather than relying on a single planning paradigm.

Introduction

Path planning is a core capability of an autonomous mobile robot (AMR) because perception and localization are useful only when the robot can transform environmental knowledge into safe motion toward a task goal. The classical formulation searches for a feasible path from an initial state to a destination while minimizing a cost such as distance, traversal time, or energy. A* established an efficient heuristic search mechanism for minimum-cost paths on graphs [1], while artificial potential fields introduced a real-time reactive interpretation in which the goal attracts the robot and obstacles generate repulsive effects [2]. These approaches

remain influential because they are simple, interpretable, and computationally practical.

Dynamic environments add a second level of difficulty. The free space changes with time, the path that was safe during global planning can become blocked a moment later, and the robot must obey velocity and acceleration limits while responding to new observations. Fox, Burgard, and Thrun addressed this by searching directly in velocity space through the Dynamic Window Approach, which evaluates admissible translational and rotational commands according to heading, clearance, and speed [3]. In sampling-based planning, RRT* established asymptotic

optimality for motion planning [4], and Informed RRT* accelerated convergence by concentrating samples inside the subset capable of improving the current solution [5]. These techniques are attractive for large continuous spaces, but a static optimal path alone does not guarantee safe execution around moving objects.

The growth of machine learning has introduced policies that learn navigation behavior from interaction. Tai, Paolo, and Liu demonstrated continuous mapless control using deep reinforcement learning [6], and Faust et al. combined probabilistic roadmaps with learned local policies in PRM-RL [7]. Predictive extensions of classical local planning are also important: Missura and Bennewitz explicitly incorporated predicted collision conditions into the DWA framework [8]. More recent pre-2021 studies show the value of hybridization. Ravankar et al. developed HPPRM by combining potential-based guidance and probabilistic roadmaps for improved dynamic planning [9], while Semnani et al. combined deep reinforcement learning with force-based motion planning for dense multi-agent environments [10]. Yu, Su, and Liao further showed that neural networks and hierarchical reinforcement learning can reduce planning time and improve path smoothness under tested scenarios [11].

These developments motivate the central methodological question of this paper: how should an AMR planner balance global route quality, local responsiveness, safety, and computation when obstacles are moving? The objective is therefore not to claim that one universal algorithm is best. Instead, the paper establishes a reproducible comparison protocol and evaluates a predictive hybrid planner against representative global and local baselines. The proposed methodology is designed for differential-drive or holonomic AMRs whose sensor data can be converted into a local occupancy representation and whose controller can accept discrete or continuous short-horizon motion commands.

Literature Review

Table 1. Representative literature used to define the methodology

Ref.	Year	Method	Main contribution	Limitation for dynamic AMRs
[1]	1968	A*	Heuristic minimum-cost graph search	Requires map updates and replanning
[2]	1986	Potential field	Real-time reactive obstacle avoidance	Can suffer local minima/oscillation
[3]	1997	DWA	Velocity-space collision avoidance	Primarily local; needs global guidance

The literature can be organized into four families. First, graph-search methods such as A* [1] are deterministic and can produce short routes on discretized maps, but frequent map changes require repeated search. Second, potential-field methods [2] offer very fast obstacle response and can be extended to moving obstacles, yet local minima and oscillations remain known concerns. Third, velocity-space local planners such as DWA [3] include robot dynamics directly and are therefore well suited to real-time collision avoidance; predictive DWA variants improve foresight when obstacle motion can be estimated [8]. Fourth, sampling and learning methods expand the planner's capability in continuous or uncertain spaces. RRT* and Informed RRT* provide systematic sampling-based optimization [4], [5], while learned policies can map high-dimensional sensor inputs to navigation actions [6], [7], [10], [11].

The strongest trend across the reviewed work is hybridization. Sampling methods are improved by informed guidance, roadmaps are paired with reinforcement-learning controllers, and reactive methods are strengthened by prediction or learned decision rules. Jaradat, Al-Rousan, and Quadan applied reinforcement learning directly to mobile-robot navigation in a dynamic environment [12], Azzabi and Nouri improved the potential-field formulation for mobile-robot path planning [13], and Lei, Zhang, and Dong used deep reinforcement learning for dynamic path planning in an unknown environment [14]. HPPRM [9] illustrates the advantage of fusing roadmap connectivity with potential-based guidance, whereas Semnani et al. [10] show that learned policies and analytical force-based logic can compensate for one another in dense interactions. This suggests that dynamic AMR navigation should be treated as a layered decision process: a global planner provides strategic direction, a prediction module estimates near-future occupancy, and a local planner produces executable collision-free motion.

[5]	2014	Informed RRT*	Focused optimal sampling after first solution	Dynamic changes may invalidate branches
[7]	2018	PRM-RL	Roadmap plus learned short-range policy	Training and roadmap dependence
[8]	2019	Predictive DWA	Anticipates future collision states	Depends on motion prediction quality
[9]	2020	HPPRM	Potential-guided probabilistic roadmap	Sampling/preprocessing overhead
[10]	2020	DRL + force planning	Hybrid dense multi-agent motion planning	Learning complexity and transfer risk
[11]	2020	NN + HRL	Learning-based path optimization and smoothness	Training and validation requirements

Methodology

The proposed methodology uses a layered global-local navigation architecture. At each control cycle, current range or occupancy information is integrated into a 2-D grid map. Static obstacles remain fixed, whereas dynamic obstacles are represented by their measured position and a short-term velocity estimate. A global planner

produces a route to the destination; a prediction layer marks cells expected to be occupied during the next one or two control steps; and a local planner selects the next safe action according to goal progress, global-path consistency, clearance, and motion smoothness. Replanning is triggered whenever the current global route intersects observed or predicted occupancy.




Figure 1. Proposed hybrid global-local methodology for intelligent path planning.

Step 1 - Environment representation. The workspace is discretized into a 30×30 occupancy grid. Each cell is either free, occupied by a fixed structure, or temporarily occupied by a moving obstacle. A robot pose is represented by $p_t = (x_t, y_t)$. For continuous robots, the same logic can be applied to a local cost map after footprint inflation. The goal cell remains fixed during an experimental episode.

Step 2 - Baseline global planning. Replanning A* uses the Euclidean distance to the goal as its heuristic and recomputes a path when dynamic occupancy changes. Its evaluation cost is $f(n) = g(n) + h(n)$, where $g(n)$ is accumulated path cost and $h(n)$ estimates remaining distance. This baseline represents a planner with strong global guidance but no explicit prediction of future obstacle motion [1].

Step 3 - Reactive baselines. The potential-field baseline assigns attraction to the goal and repulsion to nearby occupied cells, following the concept introduced by Khatib [2]. The DWA-inspired baseline evaluates neighboring motion

actions using a weighted score consisting of goal distance, obstacle clearance, motion continuity, and forward progress. This simplified implementation captures the decision structure of DWA [3] without reproducing a specific commercial robot's acceleration envelope.

Step 4 - Predictive hybrid planner. The hybrid method estimates one- and two-step future positions of moving obstacles from their recent velocity. Current and predicted occupied cells are passed to A* to form a global route. The local stage then scores admissible neighboring actions, giving the highest preference to the next global-path cell while rewarding clearance and heading continuity. Actions that intersect predicted occupancy are rejected. This design follows the principle that prediction strengthens reactive navigation [8] and that combining complementary planners can improve robustness [9], [10].

Step 5 - Replanning and termination. After the robot executes one action, dynamic obstacles update their positions and the process repeats.

An episode is successful when the robot reaches the goal without collision within 220 actions. It fails when the robot collides, no valid motion can be found for the remaining horizon, or the maximum step count is reached. The same map-generation seeds and obstacle trajectories are reused across planners to ensure paired comparison.

1. Experimental Setup and Evaluation Metrics

The benchmark contains two barrier-like structures with narrow passages, an additional horizontal barrier, randomly distributed fixed cells, and four moving obstacles. Fifty scenarios are generated using deterministic random seeds. Every planner starts at cell (1,1) and aims for (28,28). Moving obstacles travel in cardinal or diagonal directions and change direction when blocked by fixed geometry. Unlike the planner, obstacle motion does not yield to the robot; therefore, a robot that enters a cell occupied on the next time step is counted as a collision. This condition creates a genuine need for short-horizon prediction.

Five metrics are recorded: (1) success rate, the percentage of scenarios completed without collision; (2) mean path length for successful

trials, computed using unit and diagonal cell distances; (3) mean planning time per executed step in the reference implementation; (4) average obstacle clearance measured in grid-cell distance; and (5) number of direction changes, used as a discrete smoothness indicator. Planning-time values are implementation-specific and are included only for relative comparison within the benchmark. Because the purpose is methodological, no statistical claim is made about a particular commercial processor or AMR platform.

Results

The comparative results are reported in Table 2. The predictive hybrid planner completed 46 of 50 scenarios, corresponding to a 92% success rate. Replanning A* achieved 86%, while both the potential-field and DWA-inspired baselines achieved 62%. The hybrid method therefore improved completion by six percentage points over reactive A* replanning and by thirty percentage points over the two local baselines in this benchmark. Its mean successful-path length (40.93 cells) remained close to A* (40.27 cells), indicating that the extra safety logic did not cause a large global detour.

Table 2. Comparative benchmark results over 50 dynamic scenarios

Planner	Success (%)	Path length	Time/step (ms)	Avg. clearance	Direction changes
Replanning A*	86	40.27	0.27	1.87	8.44
Potential Field	62	136.21	0.05	2.27	70.52
DWA-inspired	62	78.20	0.11	2.86	25.42
Hybrid Predictive	92	40.93	0.37	1.89	8.61




Figure 2. Success rate of the evaluated planning strategies.

The potential-field method maintained relatively high average clearance but accumulated long trajectories and many direction changes, which is consistent with the tendency of purely local fields to oscillate or become trapped near competing attractive and repulsive influences [2]. The DWA-inspired baseline also kept larger average clearance, but without persistent global guidance it sometimes spent many steps making safe local progress that did not efficiently resolve the

overall route. This explains why clearance alone should not be treated as the primary performance measure.

Repeated A* produced the shortest mean route and few turns because each successful episode followed a globally directed grid path. However, its 14% failure rate illustrates the limitation of treating only current obstacle locations as blocked: a moving obstacle can enter the robot's next cell after the action is chosen. The hybrid planner avoids a portion of these conflicts by excluding predicted near-future occupancy. The extra computation increased mean planning time from 0.27 ms to 0.37 ms per step in the reference implementation, but the absolute increase was small relative to the improvement in successful completion.

The outcome supports a key methodological conclusion also visible in prior work: dynamic navigation benefits from coupling planners with

different strengths rather than optimizing a single criterion. Predictive DWA research shows the value of foresight [8], HPPRM demonstrates hybrid roadmap and potential guidance [9], and deep-RL hybrid methods demonstrate the value of switching or combining decision mechanisms in dense environments [10]. The present benchmark reaches a similar conclusion using a transparent and reproducible grid-based architecture.

Conclusion and Discussion

This methodology paper evaluated intelligent path planning for autonomous mobile robots in dynamic environments through a common benchmark and a layered planning architecture. The review showed that A*, potential fields, DWA, RRT*/Informed RRT*, and reinforcement-learning approaches solve different parts of the navigation problem rather than replacing one another. The proposed predictive hybrid method therefore uses global route information for directional consistency, local scoring for executable motion, and short-horizon obstacle prediction for collision avoidance. Across 50 randomized scenarios, the hybrid planner achieved the highest success rate while preserving near-A* path length and moderate computation. The experiment also reveals important trade-offs: large clearance does not guarantee efficient goal reaching, global optimality does not guarantee future collision avoidance, and learning-based or sampling-based intelligence still requires safety-aware execution logic. For practical AMRs, the most promising architecture is consequently hierarchical: perception and tracking estimate dynamic state, a global planner supplies strategic route information, and a local planner continuously checks kinematic feasibility and predicted collision risk. The study has limitations. The benchmark uses a 2-D grid, constant short-term obstacle velocities, and simplified local actions; it does not model sensor noise, pedestrian intent, wheel slip, communication delay, or full nonholonomic acceleration constraints. Future work should validate the methodology in ROS/Gazebo or an equivalent simulator, introduce trajectory prediction uncertainty, compare model-predictive and reinforcement-learning local controllers, and test on a physical differential-drive AMR. Nonetheless, the procedure provides a clear experimental foundation for evaluating dynamic path planning before deployment.

References

- P. E. Hart, N. J. 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. DOI: 10.1109/TSSC.1968.300136.
- O. Khatib, "Real-Time Obstacle Avoidance for Manipulators and Mobile Robots," *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986. DOI: 10.1177/027836498600500106.
- D. Fox, W. Burgard, and S. Thrun, "The Dynamic Window Approach to Collision Avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997. DOI: 10.1109/100.580977.
- S. Karaman and E. Frazzoli, "Sampling-based Algorithms for Optimal Motion Planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011. DOI: 10.1177/0278364911406761.
- J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Informed RRT*: Optimal Sampling-based Path Planning Focused via Direct Sampling of an Admissible Ellipsoidal Heuristic," *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2997–3004, 2014. DOI: 10.1109/IROS.2014.6942976.
- L. Tai, G. Paolo, and M. Liu, "Virtual-to-Real Deep Reinforcement Learning: Continuous Control of Mobile Robots for Mapless Navigation," *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 31–36, 2017. DOI: 10.1109/IROS.2017.8202134.
- Faust, O. Ramirez, M. Fiser, K. Oslund, A. Francis, J. Davidson, and L. Tapia, "PRM-RL: Long-range Robotic Navigation Tasks by Combining Reinforcement Learning and Sampling-based Planning," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5113–5120, 2018. DOI: 10.1109/ICRA.2018.8461096.
- M. Missura and M. Bennewitz, "Predictive Collision Avoidance for the Dynamic Window Approach," *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8620–8626, 2019. DOI: 10.1109/ICRA.2019.8794386.
- A. Ravankar, A. Ravankar, T. Emaru, and Y. Kobayashi, "HPPRM: Hybrid Potential Based Probabilistic Roadmap Algorithm for Improved Dynamic Path Planning of Mobile Robots," *IEEE Access*, vol. 8, pp. 221743–221766, 2020. DOI: 10.1109/ACCESS.2020.3043333.

S. H. Semnani, H. Liu, M. Everett, A. de Ruiter, and J. P. How, "Multi-Agent Motion Planning for Dense and Dynamic Environments via Deep Reinforcement Learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3221–3226, 2020. DOI: 10.1109/LRA.2020.2974695.

J. Yu, Y. Su, and Y. Liao, "The Path Planning of Mobile Robot by Neural Networks and Hierarchical Reinforcement Learning," *Frontiers in Neurorobotics*, vol. 14, article 63, 2020. DOI: 10.3389/fnbot.2020.00063.

M. A. K. Jaradat, M. Al-Rousan, and L. Quadan, "Reinforcement Based Mobile Robot Navigation in Dynamic Environment," *Robotics and*

Computer-Integrated Manufacturing, vol. 27, no. 1, pp. 135–149, 2011. DOI: 10.1016/j.rcim.2010.06.019.

Azzabi and K. Nouri, "An Advanced Potential Field Method Proposed for Mobile Robot Path Planning," *Transactions of the Institute of Measurement and Control*, vol. 41, no. 11, pp. 3132–3144, 2019. DOI: 10.1177/0142331218824393.

X. Lei, Z. Zhang, and P. Dong, "Dynamic Path Planning of Unknown Environment Based on Deep Reinforcement Learning," *Journal of Robotics*, vol. 2018, Article ID 5781591, pp. 1–10, 2018. DOI: 10.1155/2018/5781591.