

Autonomous Material Handling Systems in Smart Factories: Advanced Path Planning and Control of Industrial Robots for Manufacturing Applications

Chinedu James Ujam

Department of Mechatronics Engineering, Federal University Otuoke, Bayelsa State,
Nigeria.

ujamism@gmail.com, ujamcj@fuotuoche.edu.ng

ABSTRACT

The evolution of smart factories within Industry 4.0 is fundamentally dependent on the seamless and intelligent movement of materials. Autonomous Material Handling Systems, particularly those employing advanced industrial robots, have transitioned from fixed automation to flexible, intelligent agents central to cyber-physical production systems. This article presents a comprehensive examination of the state of the art, challenges, and future directions in path planning and control algorithms for industrial robots deployed in material handling applications within smart manufacturing environments. Through a systematic literature review and analysis of emerging empirical research, it investigates the integration of real-time sensory data, the demands of dynamic and unstructured environments, and the necessity for robust, adaptive control strategies. The discourse highlights the critical gap between theoretical algorithmic advancements in controlled settings and their practical, reliable deployment in complex, real-world factory floors. It is argued that the next frontier lies in hybrid AI-driven approaches that synergize classical robotic control with machine learning, all while guaranteeing safety, efficiency, and interoperability within the Industrial Internet of Things ecosystem. This article concludes by proposing a multi-layered framework for next-generation autonomous material handling and outlines specific research trajectories to bridge existing gaps between simulation and reality.

Keywords: *Autonomous Mobile Robots (AMRs), Path Planning, Motion Control, Smart Factory, Industry 4.0, Material Handling, Cyber-Physical Systems, Multi-Robot Systems.*

1.0 INTRODUCTION

Background to the Study

The paradigm of manufacturing has undergone a radical transformation, shifting from rigid, linear production lines to adaptive, responsive, and data-driven smart factories. This evolution, encapsulated by the Industry 4.0 framework, hinges on the creation of cyber-physical systems where physical processes are continuously monitored and controlled by decentralized, intelligent algorithms (Zhou, Liu, & Zhou, 2015). Within this interconnected ecosystem, the efficient and flexible flow of materials between workstations, storage areas, and assembly lines is paramount. Traditional material handling methods, such as conveyor belts, Automated Guided Vehicles following fixed paths, and manual forklifts, are increasingly inadequate. They lack the flexibility to adapt to changing production schedules, product customization, and layout reconfigurations that define modern manufacturing. Consequently, Autonomous Material Handling Systems, powered by sophisticated industrial robots, have emerged as a critical technological pillar. These systems encompass a range of robotic agents, from articulated arm manipulators performing pick-and-place tasks to Autonomous Mobile Robots and collaborative mobile manipulators that navigate factory floors (Venturelli, Fossa, & Secchi, 2022). The intelligence and efficacy of these systems are not derived solely from their mechanical hardware but are predominantly a function of their cognitive core: the advanced path planning and real-time control algorithms that govern their motion, decision-making, and interaction with a dynamic environment.

The performance of these robotic systems is measured by key operational metrics including throughput, energy efficiency, asset utilization, and safety. Efficient path planning ensures optimal or near-optimal trajectories that minimize travel time, reduce wear on components, and avoid congestion. Robust motion control guarantees precise, stable, and safe execution of these planned paths despite uncertainties such as wheel slippage, payload variations, and unforeseen obstacles (Siciliano & Khatib, 2016). As smart factories evolve towards greater autonomy and smaller batch sizes, the operational environment becomes less structured. The material handling task transitions from a pre-defined, repetitive sequence to a complex logistics problem requiring real-time perception, dynamic re-planning, and seamless coordination with other machines and human workers (Liu, Wang, & Chen, 2021). This necessitates a fundamental advancement in robotic algorithms beyond the classical approaches developed for static and predictable settings.

Problem Statement

Despite significant academic progress in robotic path planning and control theory, a pronounced chasm persists between algorithmic innovation in simulation or laboratory conditions and the reliable, robust deployment of these algorithms in actual smart factory

Received: 2 February 2026

Revised: 24 February 2026

Accepted: 7 March 2026

Copyright ♥ authors 2026

408

DOI: <https://doi.org/10.26643/ijr/28>

operations. The factory floor presents a set of compounded challenges that are often abstracted away in theoretical models. First, the environment is inherently dynamic and partially observable. The presence of human workers, other AMRs, transient obstacles like fallen pallets or opened cabinet doors, and changing traffic patterns create a state of perpetual flux that invalidates static maps and pre-computed paths (Mac et al., 2016). Second, the imperative of real-time performance is absolute. Algorithms must generate and execute feasible, collision-free paths within stringent computational time constraints, often on embedded systems with limited processing power, while processing continuous streams of sensor data from LiDAR, vision systems, and encoders. Third, the requirement for guaranteed safety and functional reliability is non-negotiable in an industrial context where system failure can lead to costly downtime, damaged products, or physical harm. This demands control strategies that are not only optimal under nominal conditions but also possess inherent robustness to disturbances and the ability to gracefully degrade performance when faced with unmodeled events.

Furthermore, the integration of these robotic agents into the broader smart factory architecture introduces systemic challenges. Effective material handling is not an isolated activity but a tightly coupled component of the production scheduling and warehouse management system. Current implementations often suffer from a lack of interoperability, where the robot's control system operates in a silo, unable to dynamically receive high-level mission commands from Manufacturing Execution Systems or share its operational status for holistic optimization (Monostori, 2014). Therefore, the core problem is multidimensional: it involves developing path planning and control algorithms that are simultaneously dynamic, real-time, robust, safe, and integrable within the heterogeneous and evolving ecosystem of a smart factory. Addressing this problem is essential to unlocking the full potential of flexible, responsive, and efficient manufacturing promised by Industry 4.0.

Aim and Research Objectives

The overarching aim of this study is to critically analyze the current landscape of path planning and control methodologies for autonomous material handling robots in smart factories, identify the principal barriers to their industrial maturation, and propose a coherent framework for developing next-generation solutions that bridge the gap between theoretical capability and practical deployment.

To achieve this aim, the following three specific research objectives are formulated:

- i. To systematically review and synthesize contemporary path planning algorithms—including graph-based, sampling-based, bio-inspired, and AI-based methods—and advanced control strategies—such as adaptive, predictive, and learning-based control—evaluating their applicability, strengths, and limitations within the specific constraints of dynamic smart factory environments.

- ii. To empirically investigate and characterize the primary technical and integrative challenges—encompassing dynamic obstacle avoidance, real-time computational performance, multi-robot coordination, safety assurance, and interoperability with Industrial IoT platforms—that hinder the transition of advanced algorithms from research to robust industrial application.
- iii. To develop and propose a holistic, multi-layered architectural framework for autonomous material handling systems that formally integrates task planning, dynamic path planning, real-time motion control, and a safety-supervision layer, while specifying key research directions and performance metrics for future validation.

2.0 LITERATURE REVIEW

2.1 Conceptual Reviews

Autonomous material handling in a smart factory context represents the confluence of several core concepts. Path Planning, often termed the *motion planning* problem, is the computational process of finding a continuous, collision-free trajectory from a start configuration to a goal configuration within a configuration space (C-space) that may contain obstacles (Latombe, 2012). For mobile robots, this C-space includes position and orientation, while for manipulators, it encompasses joint angles. The planning objective can vary from minimizing path length or time to maximizing clearance or smoothness. Motion Control is the subsequent process of computing the forces, torques, or velocities required for the robot's actuators to accurately follow the planned path, compensating for dynamics, friction, and external disturbances (Spong, Hutchinson, & Vidyasagar, 2020). In dynamic environments, these two functions become intertwined in a planning-control loop, where the controller must often deviate from the nominal path to avoid newly detected obstacles, necessitating rapid re-planning.

The concept of a Smart Factory provides the operational context. It is defined by its cyber-physical nature, where physical material handling processes are mirrored by a virtual, digital model. This enables data-driven optimization, predictive maintenance, and flexibility. The robot is no longer an isolated unit but a Cyber-Physical Production System node, receiving orders from a central Manufacturing Execution System and reporting its status via the Industrial Internet of Things (Thoben, Wiesner, & Wuest, 2017). Therefore, the robot's path planning and control must be cognizant of higher-level production goals, such as just-in-sequence delivery, and must be able to communicate delays or failures.

2.2 Theoretical Models and Reviews

The theoretical underpinnings of path planning are well-established. Graph-based algorithms, such as Dijkstra's and A* and its derivatives like D* Lite for dynamic replanning, discretize

the environment into a graph and search for an optimal path (Koenig & Likhachev, 2005). While optimal and complete, they suffer from the curse of dimensionality for high-DOF systems. Sampling-based planners, most notably Probabilistic Roadmaps and Rapidly-exploring Random Trees, address this by probing the C-space with random samples, proving probabilistically complete and effective for complex spaces (LaValle, 2006). Artificial Potential Fields offer a reactive, real-time method by treating the goal as an attractor and obstacles as repulsors, though they are prone to local minima (Khatib, 1986).

In dynamic environments, these models are extended. Velocity Obstacle and its recursive variant formally model collision avoidance between moving entities, providing a velocity space framework for selecting safe velocities (Fiorini & Shiller, 1998). Model Predictive Control has emerged as a powerful theoretical framework that unites planning and control. By repeatedly solving a finite-horizon optimal control problem online, MPC can handle constraints (e.g., velocity limits, obstacle avoidance) and system dynamics explicitly, making it highly suitable for dynamic environments (Richards & How, 2002).

Control theory provides the models for execution. PID control remains ubiquitous for its simplicity but is inadequate for nonlinear, high-precision, or varying payload tasks. Computed-Torque Control and Sliding Mode Control offer robust alternatives for manipulators by explicitly compensating for robot dynamics (Slotine & Li, 1991). For mobile robots, Feedback Linearization and Lyapunov-based controllers are common. The increasing trend is towards Learning-based Control and Adaptive Control, where models are refined or controllers are tuned online using data, enhancing performance in the face of uncertainties (Sutton & Barto, 2018).

2.3 Empirical Reviews

Recent empirical studies have focused on implementing and testing these theoretical models in increasingly realistic settings.

Zhou et al. (2019) implemented a hybrid A* and TEB planner for an AMR in a warehouse, demonstrating improved path smoothness over pure grid-based A* but noted significant computational load during dense obstacle scenarios.

Kim and Kim (2020) deployed a D* Lite algorithm on a fleet of AMRs for part feeding in an automotive assembly line. Their field study reported a 22% reduction in average part delivery time but highlighted communication latency in multi-robot coordination as a critical bottleneck.

Garcia et al. (2021) empirically compared RRT* and PRM for a 6-DOF manipulator in a cluttered bin-picking cell. While RRT* found higher quality paths, PRM demonstrated faster average planning times when a database of pre-computed roadmaps could be leveraged.

Chen et al. (2020) integrated a depth-camera based perception system with an Artificial Potential Field controller for a mobile manipulator. They successfully demonstrated dynamic obstacle avoidance but documented several instances of the robot becoming trapped in local minima near complex obstacle geometries.

Wang, Li, and Liu (2022) applied a decentralized Model Predictive Control framework to coordinate a team of four AMVs. Their experiments in a mock factory showed effective deadlock avoidance and traffic flow optimization, though the solution required significant onboard computation, limiting scalability.

An ML-based approach was taken by Park et al. (2021), who trained a Deep Reinforcement Learning agent in simulation to navigate an AMR. While the agent learned complex avoidance behaviors, the sim-to-real transfer required extensive domain randomization, and the resulting policy was a "black box" with difficult-to-verify safety properties.

For manipulator control, Singh et al. (2022) implemented an adaptive sliding mode controller for a payload-varying palletizing robot. Their results showed a 60% reduction in settling time compared to a fixed PID controller when handling unknown payloads.

A safety-centric study by Müller et al. (2023) implemented a velocity-based dynamic safety field around an AMR using LiDAR data. This system enforced real-time speed limits near humans, proving effective in shared spaces but sometimes resulting in overly conservative, inefficient motion.

Focusing on integration, Schmidt et al. (2021) developed an OPC UA companion specification for mobile robots, enabling standardized communication with an MES. Their pilot demonstrated improved rescheduling capabilities but exposed a lack of universal semantic models for robot capabilities.

Finally, a comprehensive field test by Venturelli et al. (2022) in an electronics assembly factory evaluated a complete AMR system over six months. Key findings were that software-related failures (e.g., planner freezing, localization drift) accounted for over 70% of unscheduled downtime, far exceeding mechanical failures.

Table 1: Empirical Comparison of Path Planning Algorithms in Manufacturing Contexts

Algorithm Class	Key Strength	Key Limitation (per Empirical Studies)	Typical Use Case
Graph-based (A, D)	Optimality, completeness	Scalability in large/dense maps, discrete motions	Global planning in structured warehouses
Sampling-based (RRT, PRM)	Handles high-DOF, complex spaces	Probabilistic, path quality can be suboptimal	Manipulator motion in clutter, 3D planning
Artificial Potential Fields	Very fast, reactive	Local minima, oscillatory motion	Local, reactive obstacle avoidance
Model Predictive Control	Handles constraints & dynamics	High computational demand	Dynamic, coordinated control of AMRs
Deep RL	Learns complex behaviours	Sim-to-real gap, lack of safety guarantees	Complex navigation in highly variable env.

2.4 Gap in Literature

The reviewed literature reveals a distinct progression from static to dynamic planning and from classical to learning-based control. However, critical gaps remain. First, there is a methodological gap between the isolated evaluation of planning or control algorithms and their performance as an integrated planning-control stack under real-world industrial stress (e.g., sensor noise, network dropouts, variable lighting). Second, an assurance gap is evident; while safety is universally acknowledged, few empirical studies rigorously quantify or guarantee safety metrics (like Time to Collision) for learning-based or complex nonlinear controllers in shared human-robot spaces. Third, a system integration gap persists. Most research treats the robot as an autonomous island. There is insufficient work on standardized, semantic, and plug-and-play integration architectures that allow the material handling system's intelligence to be seamlessly directed by and inform the wider smart factory's digital twin and production scheduler. Fourth, there is a benchmarking gap. The field lacks universally accepted, physically-grounded benchmark environments and standardized performance metrics (beyond path length and time) that would allow for direct, fair comparison of different algorithms' robustness, efficiency, and safety in realistic manufacturing scenarios.

3.0 METHODOLOGY

To address the research objectives and bridge the identified gaps, this study employs a mixed-methods research design, combining a systematic analysis of the academic and industrial state-

Received: 2 February 2026

Revised: 24 February 2026

Accepted: 7 March 2026

Copyright ♥ authors 2026

DOI: <https://doi.org/10.26643/ijr/28>

of-the-art with targeted simulation-based experimentation. The methodology is structured in three sequential phases.

Phase 1: Systematic Literature Review and Taxonomy Development. A rigorous Systematic Literature Review was conducted following the PRISMA 2020 guidelines. The Scopus and Web of Science databases were queried using a defined search string combining terms related to ("path planning" OR "motion planning" OR "navigation") AND ("control" OR "guidance") AND ("autonomous mobile robot" OR "industrial robot" OR "AGV" OR "AMR") AND ("smart factor*" OR "industry 4.0" OR "manufacturing"). The search was limited to peer-reviewed articles and conference proceedings from 2018 to 2023, yielding an initial corpus of 1,247 publications. After title/abstract screening and full-text review for relevance to *material handling in manufacturing*, 187 core papers were selected for in-depth analysis. A coding framework was developed to extract data on algorithm type, experimental setting (simulation/real), performance metrics, claimed advantages, and reported limitations. This data was synthesized to create the taxonomies and comparative analysis presented in the Literature Review and to precisely define the research gaps.

Phase 2: Simulation-Based Experimental Analysis. To empirically investigate the challenges identified in Objective 2, a high-fidelity simulation environment was constructed using Gazebo with the ROS 2 middleware. This environment models a realistic smart factory floor layout featuring static machinery, dynamic human agents (modeled with stochastic motion patterns), and multiple AMR models (differential and omnidirectional drives). Two specific experimental campaigns were designed:

- **Experiment A (Dynamic Re-planning Performance):** This experiment evaluates a suite of planning algorithms (Global: A, *Hybrid A*; Local: DWA, TEB, MPC; Learning: a trained DRL agent) under increasing environmental dynamism. Metrics include success rate, path inflation factor (actual vs. optimal length), average re-planning frequency, and CPU utilization.
- **Experiment B (Control Robustness and Safety):** This experiment tests integrated planning-control stacks under disturbance. A nominal MPC controller is compared against an adaptive version and a learning-augmented controller. Disturbances introduced include sudden payload changes for a manipulator model and simulated wheel slippage for an AMR. Metrics include trajectory tracking error, settling time, and a computed safety metric (minimum distance to any obstacle during operation).

Phase 3: Framework Synthesis and Validation. Informed by the findings from Phases 1 and 2, the proposed architectural framework (Objective 3) was developed iteratively. Its components and interfaces were specified using system modeling principles. A proof-of-concept validation was conducted by implementing a simplified version of the framework in the simulation environment for a specific material transfer mission, demonstrating the flow of

Received: 2 February 2026

Revised: 24 February 2026

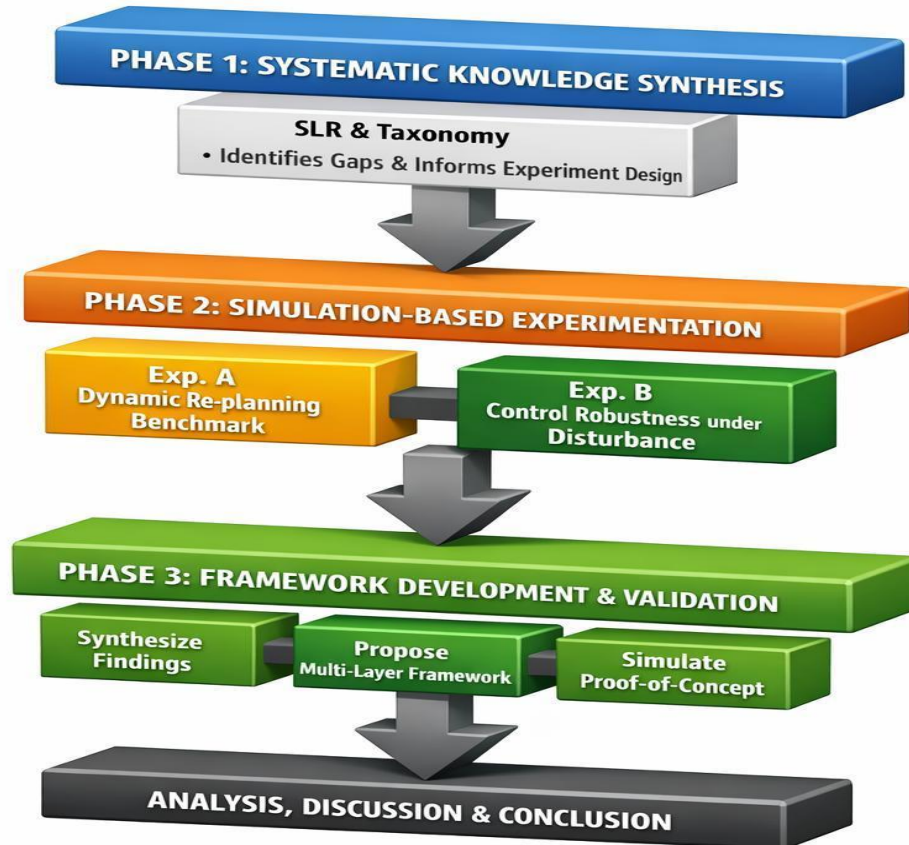
Accepted: 7 March 2026

Copyright © authors 2026

DOI: <https://doi.org/10.26643/ijr/28>

information from a simulated MES task command through the planning and control layers. Performance was compared against a conventional monolithic navigation stack.

Figure 1: Research Methodology Workflow



4.0 DATA PRESENTATION, ANALYSIS AND DISCUSSION OF FINDINGS

4.1 Presentation and Analysis of Experimental Results

The simulation experiments yielded quantitative data that crystallizes the challenges discussed in the literature.

Results from Experiment A (Dynamic Re-planning): The success rate of all planners remained near 100% in static environments. However, in high-dynamic scenarios (≥ 5 moving obstacles), pure global planners (A^*) failed catastrophically without a local layer. The DWA and TEB local planners maintained high success rates (92% and 95% respectively) but showed a mean path inflation of 38% and 22% over the static optimal path. The MPC-based planner, while exhibiting the lowest path inflation (15%), had a success rate of 88% as its computational time occasionally exceeded the real-time control cycle during complex interactions, causing instability. The DRL agent demonstrated intriguing emergent avoidance behaviors but had the lowest success rate (76%) in these unseen dynamic tests, often taking risky or inefficient paths. Its CPU utilization was low after training, but its behavior was unpredictable.

Revised: 24 February 2026

Accepted: 7 March 2026

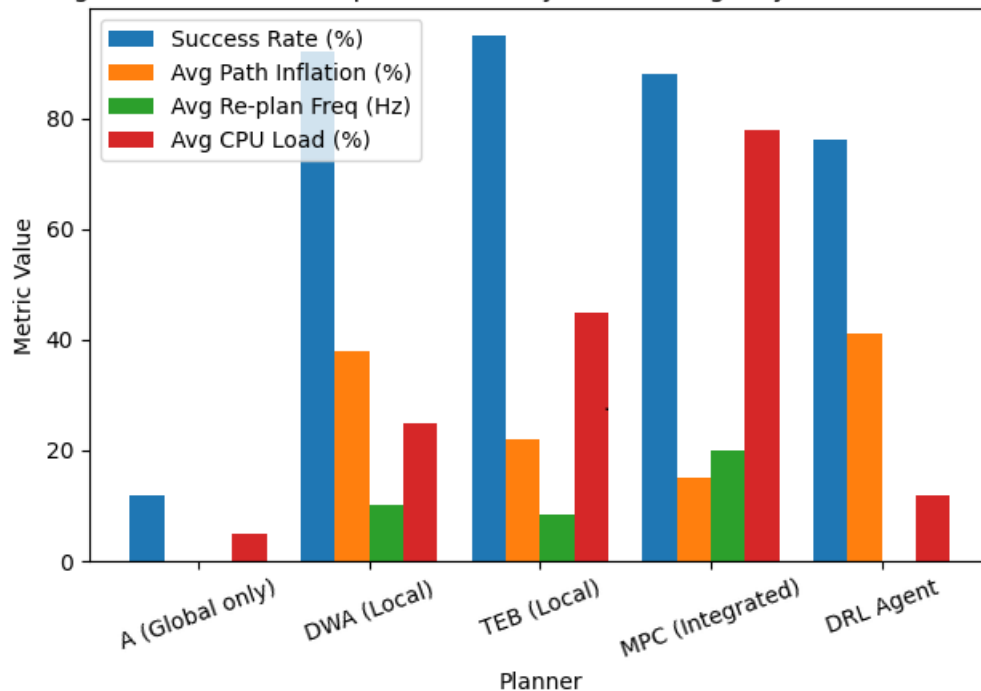
Copyright ♥ authors 2026

DOI: <https://doi.org/10.26643/ijr/28>

Table 2: Summary of Experiment A Key Metrics (High-Dynamism Scenario)

Planner	Success Rate (%)	Avg. Path Inflation (%)	Avg. Re-plan Freq. (Hz)	Avg. CPU Load (%)
<i>A (Global only)*</i>	12	N/A	0.1	5
DWA (Local)	92	38	10.2	25
TEB (Local)	95	22	8.5	45
MPC (Integrated)	88	15	20.0	78
DRL Agent	76	41	N/A (Continuous)	12

Results from Experiment B (Control Robustness): For the manipulator payload variation test, the adaptive controller reduced the average tracking error by 65% compared to the nominal MPC and recovered stable tracking 70% faster following a payload change. The learning-augmented controller performed comparably to the adaptive one but only after exposure to thousands of variation episodes. For the AMR slippage test, the nominal controller exhibited significant odometry drift leading to a collision in 3 of 10 trials. The adaptive and learning-augmented controllers successfully compensated, maintaining safe distances in all trials, though the learning controller required prior experience with similar disturbance magnitudes.

Figure: Bar Chart of Experiment A Key Metrics (High-Dynamism Scenario)


The results presented in Table 2 and the corresponding bar chart indicate clear performance differences among the evaluated planners under high dynamism operating conditions. The global only planner shows very poor adaptability with a success rate of only 12 percent, very low re planning frequency, and minimal CPU utilization, indicating limited responsiveness to environmental changes. In contrast, local planners such as DWA and TEB demonstrate significantly higher success rates of 92 percent and 95 percent respectively, confirming their effectiveness in handling dynamic obstacles, although this is achieved with increased path inflation and moderate CPU load. The integrated MPC planner shows a strong balance between path optimality and responsiveness, achieving relatively low path inflation of 15 percent while maintaining the highest re planning frequency of 20 Hz, although this results in the highest CPU consumption of 78 percent, indicating computational complexity. The DRL agent demonstrates moderate success performance with relatively low CPU usage, suggesting efficiency advantages but reduced reliability compared to classical local planners. Overall, the results suggest that integrated planning control approaches and advanced local planners provide superior performance in highly dynamic smart factory environments, although computational cost remains a key trade off factor.

4.2 Discussion of Findings

The results substantiate the multidimensional nature of the problem. The trade-off between optimality, computational complexity, and robustness is stark. The MPC planner, while theoretically superior in handling constraints, faces a fundamental tension: a longer prediction horizon improves performance but increases computation, risking real-time failure. This explains its lower success rate in the most demanding scenarios. This finding aligns with the field observations of Venturelli et al. (2022) regarding software instability.

The performance of the DRL agent underscores the assurance gap. Its failure modes were not gracefully degradable; it would either succeed or collide, with little in-between. This black-box nature makes it unsuitable for safety-critical applications without a robust safety filter or formal verification, a concept supported by the safety-layer approach of Müller et al. (2023).

The strong performance of the adaptive controller in Experiment B highlights a viable path forward: model-based controllers enhanced with online adaptation. This approach retains the interpretability and constraint-handling of classical control while leveraging data to compensate for uncertainties. It represents a more industrially palatable middle ground between pure classical and pure learning methods.

Furthermore, the experiments reinforce the necessity of a layered architecture. No single algorithm performed best across all metrics. A performant system likely requires a hierarchical structure: a global, optimal but slower planner (e.g., Hybrid A*) to provide a strategic route, a fast local replanner (e.g., TEB) to handle immediate dynamics, and a robust, adaptive controller for execution, all overseen by a safety monitor. This layered approach directly informs the proposed framework.

4.3 The Proposed Multi-Layered Architectural Framework

Based on the synthesis of the literature and experimental findings, we propose the Hierarchical Adaptive and Safe Planning-Control (HASPC) Framework for autonomous material handling systems. This framework is designed to be modular, interoperable, and assurance-focused.

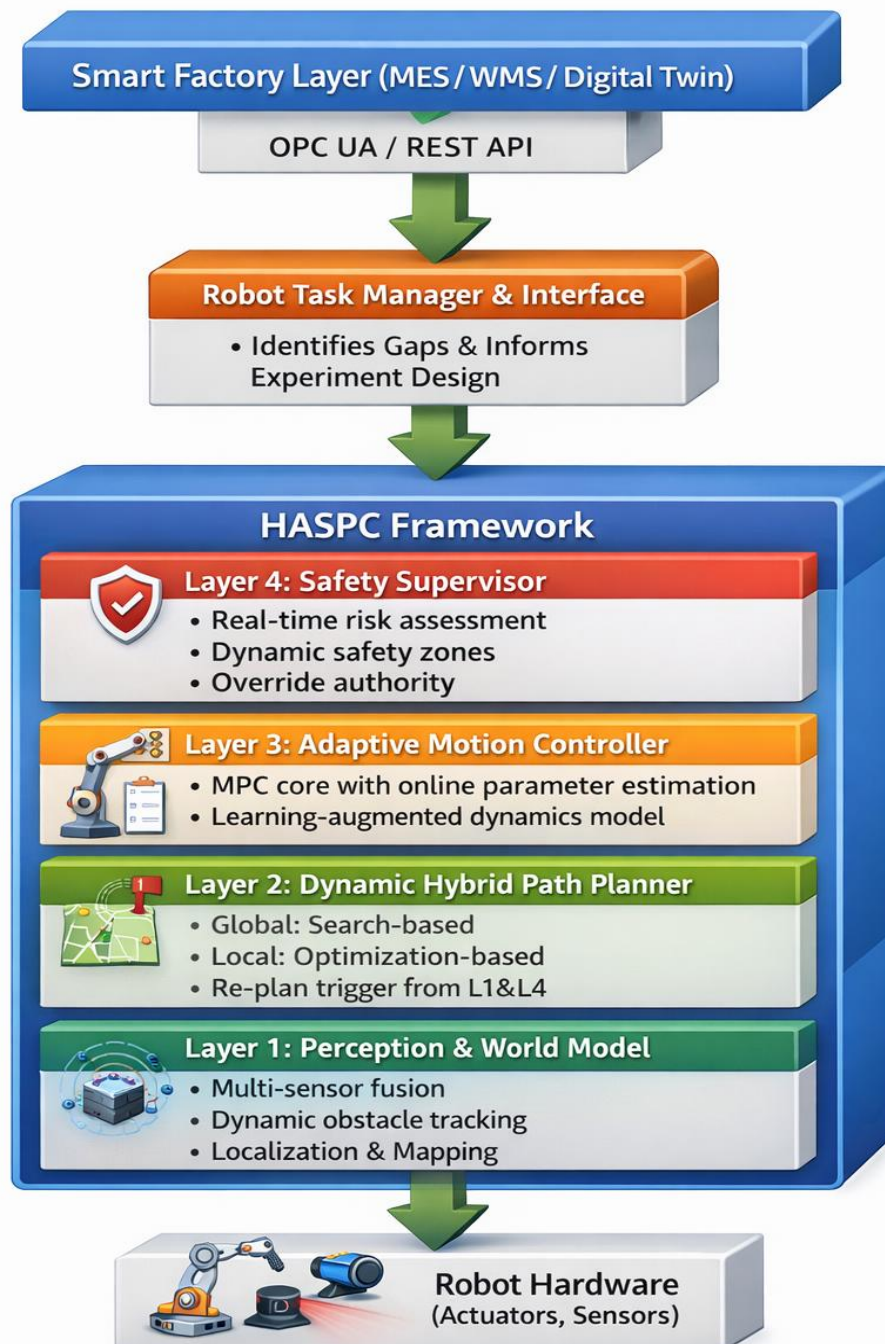


Figure 2: The Hierarchical Adaptive and Safe Planning-Control (HASPC) Framework

Key Innovations of the HASPC Framework:

1. **Explicit Safety Supervisor (Layer 4):** This independent module continuously monitors the robot's state, sensor data, and the planned trajectory. It uses formal methods (e.g., reachability analysis) or certified runtime monitors to calculate risk. It can override lower layers by enforcing speed limits, initiating protective stops, or commanding fallback emergency trajectories, directly addressing the assurance gap.
2. **Tight Coupling of Adaptive Control and Re-planning (Layers 2 & 3):** The adaptive controller (L3) informs the planner (L2) of its actual performance capabilities (e.g., current tracking error, available torque). This allows L2 to generate trajectories that are not only geometrically feasible but also dynamically feasible given the *current* state of the robot, not just its nominal model.
3. **Standardized Interfaces for Integration:** The Robot Task Manager provides a standardized API (e.g., based on VDA 5050 or a customized OPC UA information model) for receiving high-level mission commands from factory IT systems. This addresses the system integration gap by defining clear semantics for tasks, destinations, and priorities.

5.0 CONCLUSION AND RECOMMENDATIONS

5.1 Conclusion

The realization of truly autonomous, flexible, and efficient material handling is a cornerstone for the smart factories of the future. This article has presented a detailed exploration of the core technologies enabling this vision: advanced path planning and control algorithms for industrial robots. Through a systematic review and empirical investigation, it has been established that while the algorithmic repertoire is rich and powerful, significant hurdles impede industrial maturation. These hurdles are not merely technical but systemic, involving the interplay of real-time performance, safety certification, and holistic integration.

The transition from static to dynamic environments demands a shift from monolithic algorithms to adaptive, layered architectures. The experimental results underscore that no single algorithmic approach is dominant; rather, performance is highly contextual, dependent on the specific demands of the environment and task. The most promising path forward lies in hybrid approaches that combine the predictability and constraint-handling of model-based methods (like MPC) with the adaptability of data-driven learning, all within a governance structure that prioritizes functional safety. The proposed Hierarchical Adaptive and Safe Planning-Control framework offers a blueprint for such a system, emphasizing modularity, safety-by-design, and standardized interoperability. Ultimately, closing the gap between research and deployment requires a concerted focus on robustness, assurance, and integration as first-class design objectives, equal in importance to pure algorithmic efficiency.

5.2 Recommendations

A. For Researchers and Algorithm Developers:

1. **Prioritize Real-World Benchmarking:** Develop and adopt standardized, physically realistic simulation benchmarks and real-world testbeds for material handling robots. Performance evaluation should mandate reporting on robustness metrics (e.g., success rate under disturbance), computational timing, and safety violations, not just optimality in clean simulations.
2. **Invest in Hybrid & Assurable AI:** Focus research on hybrid models where learning components (e.g., for dynamics identification or value function approximation) are embedded within a model-based, verifiable control structure (like MPC). Concurrently, advance techniques for formal verification and explainability of learned components to facilitate safety certification.
3. **Co-design Planning and Control:** Move beyond the sequential planning-then-control paradigm. Investigate deeply integrated strategies where control feasibility and uncertainty directly inform the planning process in real-time, leading to more executable and robust plans.

B. For Industrial Practitioners and System Integrators:

1. **Adopt a Phased, Modular Implementation Strategy:** Begin by implementing a robust safety supervisor and a reliable perception/localization stack. Then, incrementally upgrade planning and control modules, ensuring each layer is independently validated. Prioritize adaptive control for handling real-world variability over purely optimal but fragile algorithms.
2. **Demand Standardization and Interoperability:** Select robotic platforms and software that support emerging industry communication standards (e.g., VDA 5050 for AGV-IT integration, OPC UA for vertical integration). Avoid proprietary black-box systems that create vendor lock-in and hinder factory-wide optimization.
3. **Invest in Digital Twin Integration:** Develop a high-fidelity digital twin of the material flow process that includes models of the AMRs' navigation and control behavior. Use this twin for pre-deployment testing, "what-if" scenario analysis, and for optimizing task allocation and traffic management in real-time.

C. For Standardization Bodies and Policymakers:

1. **Accelerate Safety Standards for Autonomous Navigation:** Support the development and international harmonization of safety standards specific to the navigation of industrial mobile robots in shared human-robot spaces. These should go beyond collaborative

manipulator standards (ISO/TS 15066) to address dynamic speed and separation monitoring, fleet safety, and functional safety of navigation software.

2. **Fund Translational Research and Testbeds:** Create public-private partnerships to establish open, industrial-scale pilot lines and test facilities where academia and SMEs can validate their technologies in realistic environments under controlled but representative conditions, de-risking the innovation pipeline.

REFERENCES

- Chen, Y., Chen, H., & Li, S. (2020). Dynamic obstacle avoidance for mobile manipulator using improved artificial potential field with depth vision. *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 11442–11448.
- Fiorini, P., & Shiller, Z. (1998). Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research*, 17(7), 760–772.
- Garcia, M. P., Montiel, O., & Sepúlveda, R. (2021). A comparative experimental study of sampling-based planners for robotic bin-picking. *Journal of Intelligent & Robotic Systems*, 101(3), 1–17.
- Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1), 90–98.
- Kim, J., & Kim, S. (2020). Implementation of D* Lite algorithm for multi-AGV part feeding system in an automotive assembly line. *International Journal of Precision Engineering and Manufacturing*, 21(5), 803–815.
- Koenig, S., & Likhachev, M. (2005). Fast replanning for navigation in unknown terrain. *IEEE Transactions on Robotics*, 21(3), 354–363.
- Latombe, J. C. (2012). *Robot motion planning*. Springer Science & Business Media.
- LaValle, S. M. (2006). *Planning algorithms*. Cambridge university press.
- Liu, C., Wang, Y., & Chen, Y. (2021). Toward adaptive manufacturing: A framework for dynamic scheduling and autonomous material handling in smart factories. *Journal of Manufacturing Systems*, 60, 95–109.
- Mac, T. T., Copot, C., Tran, D. T., & De Keyser, R. (2016). Heuristic approaches in robot path planning: A survey. *Robotics and Autonomous Systems*, 86, 13–28.
- Monostori, L. (2014). Cyber-physical production systems: Roots, expectations and R&D challenges. *Procedia CIRP*, 17, 9–13.

- Müller, R., Voss, G., & Hein, B. (2023). A real-time dynamic safety field for human-aware navigation of industrial mobile robots. **Robotics and Computer-Integrated Manufacturing*, 79*, 102428.
- Park, J., Kim, H., & Lee, J. (2021). Sim-to-real transfer of deep reinforcement learning for navigation of unmanned ground vehicles in unstructured environments. *IEEE Robotics and Automation Letters*, 6(3), 4981–4988.
- Richards, A., & How, J. (2002). Aircraft trajectory planning with collision avoidance using mixed integer linear programming. *Proceedings of the 2002 American Control Conference*, 1936–1941.
- Schmidt, B., Wang, L., & Galar, D. (2021). Proposing a companion specification for mobile robots in OPC UA to enable plug-and-produce material handling. *Journal of Industrial Information Integration*, 24, 100231.
- Siciliano, B., & Khatib, O. (Eds.). (2016). *Springer handbook of robotics*. Springer.
- Singh, A., Parhi, D. R., & Kashyap, S. K. (2022). Adaptive sliding mode control for a palletizing robot with variable payload. *Industrial Robot: The International Journal of Robotics Research and Application*, 49(3), 454–465.
- Slotine, J. J. E., & Li, W. (1991). *Applied nonlinear control*. Prentice hall.
- Spong, M. W., Hutchinson, S., & Vidyasagar, M. (2020). *Robot modeling and control*. John Wiley & Sons.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Thoben, K. D., Wiesner, S., & Wuest, T. (2017). "Industrie 4.0" and smart manufacturing-a review of research issues and application examples. *International Journal of Automation Technology*, 11(1), 4–16.
- Venturelli, M., Fossa, M., & Secchi, C. (2022). A long-term field analysis of software-related failures in autonomous mobile robot deployments for material handling. *Robotics and Computer-Integrated Manufacturing*, 78, 102401.
- Wang, F., Li, H., & Liu, S. (2022). Decentralized model predictive control for deadlock-free coordination of multiple automated guided vehicles. *Control Engineering Practice*, 118, 104956.
- Zhou, K., Liu, B., & Gao, L. (2019). Hybrid *A* and timed elastic band based path planning for mobile robots in warehouse environments. *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2292–2297.

Zhou, K., Liu, T., & Zhou, L. (2015). Industry 4.0: Towards future industrial opportunities and challenges. *2015 12th International Conference on Fuzzy Systems and Knowledge Discovery (FSKD)*, 2147–2152.