

AI-Driven Operational Improvement in Warehouse Management: An Integrated Framework for Demand Forecasting, Autonomous Fulfillment, Layout Optimization, and Profitability

Arvind Badrinarayanan

Spire Integrated Solutions, USA

Abstract: Warehouse operations represent a critical profitability lever in modern supply chains, yet most organizations remain dependent on reactive, siloed management approaches that cannot accommodate the speed, accuracy, and intelligence requirements of contemporary e-commerce and omnichannel fulfillment. This article proposes an integrated four-pillar artificial intelligence (AI) framework for warehouse operational improvement, structured around demand forecasting and inventory optimization, autonomous robotic order fulfillment, AI-powered slotting and layout optimization, and real-time operational intelligence. Based on peer-reviewed evidence and documented industrial deployments, the framework shows that AI-driven demand forecasting reduces forecast error by 20–40% compared to traditional autoregressive integrated moving average (ARIMA) baselines; autonomous picking systems achieve throughput of 300–600 picks per hour, representing a 3–5× improvement over manual operations; AI slotting optimization reduces picker travel time by 20–35%; and real-time AI dashboards reduce order fulfillment cycle times by 25–35%. The four pillars operate synergistically: improvements in demand accuracy reduce inventory holding costs, autonomous fulfillment reduces labor cost per unit, layout optimization increases throughput capacity, and operational intelligence converts reactive firefighting into proactive performance management. A phased implementation pathway aligned with organizational data maturity and investment capacity is presented, making AI-driven warehouse transformation accessible beyond large enterprise operators to encompass small and medium distribution enterprises. The framework advances a profitability-linked view of AI adoption, positioning warehouse intelligence as a strategic competitive differentiator rather than a cost-containment initiative.

Keywords: Artificial Intelligence, Warehouse Management, Demand Forecasting, Autonomous Mobile Robots, Inventory Optimization, Order Fulfillment, Layout Optimization, Industry 4.0

1. Introduction

Warehouse operations have emerged as one of the most strategically significant cost and service drivers in global supply chains. The global warehousing market was valued at approximately \$1.6 trillion in 2023 and continues to expand as e-commerce penetration accelerates demand for faster, more precise fulfillment. Customer expectations have undergone a structural shift: same-day and next-day delivery—once premium service tiers—are rapidly becoming standard requirements in business-to-consumer (B2C) and increasingly in business-to-business (B2B) fulfillment contexts. Amazon has demonstrated at industrial scale that AI-driven inventory placement and autonomous fulfillment are the operational foundations of same-day delivery economics at catalog breadths that manual management cannot sustain [3]. The competitive consequences for organizations failing to close the AI capability gap are concrete: higher fulfillment costs, slower delivery times, and the customer attrition that follows service quality deterioration.



Despite this urgency, the majority of warehouse operations continue to depend on reactive, siloed management approaches built around legacy warehouse management systems (WMS) that were not designed for real-time AI integration. Traditional WMS platforms provide transactional records—inventory positions, order statuses, labor activities—but lack the predictive and adaptive intelligence needed to anticipate demand shifts, optimize labor allocation dynamically, or learn from operational patterns to reduce inefficiency over time. The result is a persistent and costly operational gap: stockouts and overstock situations that tie up working capital and disrupt fulfillment, picking operations that account for 55–65% of total warehouse labor cost yet are managed through static routing rules, and layout configurations that are updated infrequently and cannot adapt to the velocity changes that seasonal demand and product portfolio evolution continuously introduce [2][4].

Artificial intelligence offers a fundamentally different operational model for warehouse management: one that is predictive rather than reactive, adaptive rather than static, and continuously learning rather than rule-bound. Machine learning-based demand forecasting replaces periodic statistical projections with continuously updated, multi-signal predictive models. Autonomous mobile robots (AMRs) and AI-powered goods-to-person systems transform order fulfillment from a labor-intensive manual process to a high-throughput, high-accuracy automated operation [3]. Dynamic slotting algorithms reposition inventory in real time based on current demand velocity. Operational intelligence dashboards surface anomalies and fulfillment risks before they breach service level agreements (SLAs), enabling managers to act on predictions rather than problems [6].

The research literature has addressed these AI applications in warehouse contexts predominantly as separate domains—forecasting systems, robotics, slotting algorithms, and dashboards treated as independent technical contributions rather than as components of an integrated operational architecture. What is absent is a unified, profitability-linked framework that positions these four capabilities as synergistic pillars of a coherent transformation strategy, with each pillar amplifying the value delivered by the others. This gap limits practitioner ability to plan coordinated AI adoption and leaves the compounding profitability effects of integrated AI deployment underexplored [1][2].

The primary contributions of this article are threefold: (1) an integrated four-pillar AI warehouse operational framework linking demand forecasting, autonomous fulfillment, layout optimization, and real-time intelligence to measurable profitability outcomes; (2) a synthesis of peer-reviewed quantitative performance evidence for each pillar; and (3) a phased implementation pathway that enables organizations at varying levels of data maturity and investment capacity to sequence their AI adoption effectively. The framework is designed to be applicable across warehouse scales and sectors, from large automated distribution centers to small and medium enterprises (SMEs) implementing their first AI capabilities.

The remainder of this article is structured as follows. Section 2 characterizes the core operational challenges that AI addresses in warehouse management. Sections 3 through 6 present each of the four framework pillars with supporting evidence. Section 7 describes the phased implementation pathway and addresses key challenges. Section 8 concludes with a synthesis and directions for future research.

2. Warehouse Management Challenges and the AI Opportunity

Demand volatility represents the foundational operational challenge in warehouse management, cascading its effects across every downstream function from inventory positioning to labor planning and transportation scheduling. Traditional demand forecasting methods—autoregressive integrated moving average (ARIMA), exponential smoothing, and their seasonal variants—generate point forecasts from historical sales data in isolation, failing to incorporate the external signals and nonlinear interaction effects that drive actual demand patterns in dynamic retail and distribution environments. The consequences of forecast inaccuracy are directly measurable: stockouts cost organizations an estimated 4–8% of annual revenue in lost sales and customer defection, while overstock situations tie up 20–30% of working capital in excess inventory that consumes storage space, generates obsolescence risk, and reduces warehouse throughput capacity [1].

Order fulfillment inefficiency compounds the financial impact of poor demand planning. Picking—the process of selecting individual items from storage locations to fulfill customer orders—accounts for 55–65% of total warehouse operating cost and remains the most labor-intensive, time-sensitive, and error-prone function in most operations [4]. Manual picking operations generate order error rates of 1–3%, which translate directly into returns processing costs, customer service escalations, and reputational damage in high-frequency consumer fulfillment environments. The fundamental constraint of manual picking—that throughput scales linearly with labor investment and is bounded by human physical capacity—creates a structural ceiling on fulfillment performance that is inconsistent with the throughput demands of high-volume e-commerce operations [3][5].

Static warehouse layout and slotting configurations amplify inefficiency by forcing workers and automated systems to travel suboptimal paths through storage areas. In most warehouse operations, slot assignments are determined through periodic manual analysis and remain fixed for extended periods—months or years—while demand velocity distributions shift continuously. When fast-moving items are located in distant storage zones while slow-moving items occupy prime pick-face positions, picker travel time inflates proportionally with the mismatch between slot assignments and actual velocity. Research consistently identifies picker travel as accounting for 50–60% of total picking time, making slotting optimization one of the highest-leverage interventions available for reducing labor cost and increasing throughput without capital investment [7].

Data fragmentation across siloed WMS, enterprise resource planning (ERP), and transportation management system (TMS) platforms prevents the real-time operational visibility that AI-driven warehouse management requires. Decisions made on data that is hours or days stale—common in environments where systems synchronize in batch cycles—result in reactive management: operators discovering stockouts after they have caused fulfillment failures, identifying bottlenecks after they have disrupted throughput, and escalating SLA breaches after they have disappointed customers. The integration of these data streams into a unified operational intelligence layer is both a technical prerequisite and a strategic imperative for AI-driven warehouse transformation [6].

3. AI-Driven Demand Forecasting and Inventory Optimization

Machine learning-based demand forecasting represents a structural departure from the statistical paradigm that has governed supply chain planning for decades. Where ARIMA and exponential smoothing models extract signal from univariate sales history and apply parametric assumptions about trend and seasonality, machine learning models—including long short-term memory (LSTM) recurrent neural networks, gradient boosting frameworks, and ensemble architectures—learn complex nonlinear relationships across multi-dimensional feature spaces incorporating sales history, promotional calendars, price elasticity signals, weather patterns, macroeconomic indicators, and real-time point-of-sale (POS) feeds [1]. This multi-signal learning capability enables demand sensing at temporal granularities and accuracy levels that statistical methods structurally cannot achieve. Walter et al.'s systematic review of AI applications in supply chain demand planning documents consistent improvements of 20–40% in forecast error metrics—MAPE, RMSE, and WAPE—across retail, distribution, and manufacturing deployment contexts [1].

Dynamic safety stock and reorder point optimization is a high-value downstream application of improved forecast accuracy. Traditional safety stock formulas—derived from assumed normal demand distributions with fixed service level parameters—generate static inventory buffers that are miscalibrated across the demand distribution: excessive for stable demand items and insufficient for volatile ones. Machine learning models replace these static parameters with continuously updated, demand-signal-responsive inventory policies that shrink safety stock for predictable items while expanding buffers ahead of detected demand surges [6]. Intelligent warehousing systems combining XGBoost and Random Forest models with IoT sensor feeds have demonstrated a 32% reduction in forecast mean absolute error compared to ARIMA baselines and a 35% reduction in operational costs through predictive inventory optimization [6].

Amazon's mastery of AI-driven inventory placement across its fulfillment network provides the highest-profile industrial validation of demand forecasting intelligence at scale. By decomposing demand signals at the individual SKU-by-geography level and positioning inventory proximate to anticipated demand concentrations across its

fulfillment network nodes, Amazon has achieved same-day delivery capability for large fractions of its catalog—an outcome that is structurally impossible without the ML-based demand forecasting that enables pre-positioning decisions weeks before orders are placed [3]. This inventory positioning intelligence translates directly into reduced outbound transportation cost, shorter customer delivery windows, and higher order fill rates—the three primary financial levers connecting fulfillment operations to business profitability.

Human oversight remains essential in AI-augmented demand planning. Machine learning models perform well within the distributions represented in their training data but degrade under conditions they have not encountered: new product introductions, unprecedented promotional events, supply disruptions, and macroeconomic regime shifts that break historical demand patterns. Demand planners exercising domain judgment—injecting commercial intelligence about upcoming launches, adjusting model outputs for known market events, and supervising AI recommendations before they trigger procurement actions—maintain the quality assurance layer that prevents miscalibrated AI forecasts from generating costly inventory errors [1][4].

4. Autonomous Robotic Order Fulfillment

Autonomous mobile robots (AMRs) and goods-to-person fulfillment systems represent the most transformative AI application in warehouse operations, replacing the physical walking that dominates manual picking time with robotic transport while concentrating human effort on the cognitive and dexterous picking tasks that automation cannot yet reliably perform. AMR-based picking systems achieve throughput rates of 300–600 picks per hour—a 3–5× improvement over the 80–120 picks per hour achievable through manual operations—by eliminating the picker travel that constitutes 50–60% of total manual picking time [3][5]. Order accuracy in AMR-assisted environments improves from the 97–99% range of careful manual operations to 99.8–99.9% accuracy achieved through barcode verification at the pick face, eliminating the costly return and reprocessing workflows that accompany manual error rates.

AI-driven order batching and routing algorithms are the intelligence layer that determines the efficiency of robotic fulfillment systems. Multi-order wave batching—grouping individual customer orders whose pick paths can be efficiently combined into a single robot tour—reduces the number of robot journeys required per order, increasing system throughput without proportional robot fleet expansion. Dynamic path planning algorithms minimize robot congestion in high-density storage environments by coordinating fleet movements across shared aisle networks in real time, preventing the throughput degradation that uncoordinated robot traffic would otherwise cause [5]. Mahmoudinazlou et al.'s deep reinforcement learning (DRL) framework for dynamic order picking demonstrates that AI-based routing policies significantly outperform static routing heuristics, achieving greater throughput under variable order arrival patterns that reflect real-world demand dynamics [5].

Human-robot collaboration in fulfillment operations preserves the role of warehouse workers while elevating their productivity and reducing ergonomic strain. In goods-to-person configurations, robots transport storage pods to stationary pick stations where human pickers select items and place them into order containers—eliminating worker travel while retaining human dexterity for the varied item selection tasks that robotic grasping systems cannot yet handle reliably across diverse product categories [7]. In picker-to-parts AMR-assisted configurations, robots carry order totes alongside walking pickers, eliminating the need for pickers to push trolleys while simultaneously processing the routing optimization that determines the most efficient pick sequence [3]. Zhen et al. demonstrate that optimal AMR scheduling in picker-to-parts systems reduces picker walking distance by 30–50% while improving overall system throughput by 20–35% compared to unassisted manual picking [3].

The measurable fulfillment performance improvements enabled by robotic AI systems translate directly into competitive operational outcomes. Order fulfillment cycle time reductions of 25–35% enable faster customer delivery windows without expanded labor investment. Labor cost per unit shipped reductions of 20–40% improve contribution margins on fulfillment operations. Fakhrai Rad et al.'s empirical study of an AI-based order picking system confirms that AI-driven picking reduces travel distance and time significantly, while identifying critical success factors—including system integration quality, workforce training adequacy, and human-robot task allocation design—that determine whether the full performance potential of AI picking is realized in deployment [4].

5. AI-Powered Warehouse Layout Optimization and Slotting Intelligence

Dynamic slotting intelligence—the continuous AI-driven repositioning of SKUs to storage locations aligned with current demand velocity—is one of the highest-leverage, lowest-capital AI interventions available to warehouse operators. Traditional slotting relies on periodic manual analysis that assigns items to storage zones based on ABC velocity classification and ergonomic criteria and then remains static until the next review cycle. In practice, this means that slot assignments drift out of alignment with actual demand velocity continuously, with fast-moving seasonal items remaining in secondary locations while slower items occupy premium golden zone positions near dispatch. AI velocity analysis models—operating on rolling transactional data with daily or weekly update cycles—identify these misalignments and generate optimized slot recommendations that reposition items before the mismatch inflates travel time [7].

Space utilization optimization extends AI layout intelligence beyond pick face positioning to the full volumetric assignment of inventory across storage zones. AI-driven storage assignment algorithms balance multiple competing constraints simultaneously: pick accessibility for high-velocity items, space efficiency for bulky slow movers, product compatibility constraints for hazardous or temperature-sensitive items, and replenishment logistics that minimize restocking labor. By solving these multi-constraint assignment problems continuously across the full SKU catalog rather than periodically for subsets, AI systems identify storage configurations that improve overall warehouse capacity utilization by 15–25% without physical expansion—deferring capital investment in facility extension while supporting business growth [2].

Digital twin simulation represents the AI-powered planning capability that enables warehouse managers to evaluate alternative layout configurations before committing to physical changes. By building a virtual replica of warehouse operations—incorporating storage layout, robot fleet configurations, pick path algorithms, and labor allocation policies—digital twin platforms allow planners to simulate the throughput, labor cost, and SLA performance implications of proposed changes under a range of demand scenarios before implementation. This simulation-based planning replaces the empirical trial-and-error that has historically governed warehouse layout decisions with data-driven investment analysis, reducing the risk of costly layout changes that fail to deliver anticipated productivity improvements [6].

The continuous optimization architecture of AI slotting systems transforms layout management from a periodic project into an ongoing operational capability. As new products are introduced, seasonal demand patterns shift, and fulfillment velocity distributions evolve, AI slotting models update their recommendations on rolling cycles—daily, weekly, or event-triggered—ensuring that slot assignments track current operational reality rather than historical snapshots. This continuous adaptation is particularly valuable in high-mix fulfillment environments where SKU portfolios expand and contract frequently, and where manual slotting review cycles would require disproportionate analyst time to maintain currency [7].

6. Real-Time Operational Intelligence and Performance Management

Real-time AI operational dashboards transform warehouse management from a backward-looking performance reporting exercise into a forward-looking performance management discipline. By ingesting continuous data streams from WMS, AMR fleet management systems, labor tracking platforms, and order management systems, AI analytics layers construct a live operational model of warehouse state—tracking throughput rates, labor productivity, inventory accuracy, order queue depth, and dock scheduling—and applying anomaly detection algorithms to identify deviations from expected performance trajectories before they accumulate into SLA breaches or shift-level throughput shortfalls [6]. Machine learning forecasting of picking cycle time—as demonstrated by the real-time decision-making system for cycle time prediction developed by *Unlocking Real-Time Decision-Making in Warehouses (2024)*—enables warehouse supervisors to identify individual orders at risk of missing SLA windows hours before dispatch deadlines, supporting proactive resequencing that protects service commitments without reactive expediting costs.

Predictive fulfillment SLA management converts order-level AI forecasts into actionable operational interventions. When AI models flag an order as high-risk for late dispatch—based on current queue depth, available labor, in-process pick completion rates, and dock scheduling—supervisors receive prioritized alerts that enable targeted intervention: reallocating pickers to high-priority orders, triggering AMR routing reprioritization for at-risk wave groups, or initiating proactive customer communication before the SLA breach materializes. This prediction-driven management model reduces the operational waste and customer service cost of reactive SLA breach management, where recovery actions are typically more expensive and less effective than prevention [5][6].

The integration of operational metrics with financial performance reporting creates the closed-loop profitability management architecture that is the highest-maturity form of AI warehouse intelligence. Traditional warehouse KPI reporting—inventory accuracy, pick rate, on-time shipment rate—measures operational activity but does not directly connect operational decisions to profit and loss (P&L) outcomes. AI-enabled profitability dashboards link warehouse operational variables to financial outcomes in real time: inventory carrying cost as a function of current stock positions and forecast accuracy; labor cost per unit as a function of pick rate, batch efficiency, and robotic utilization; and revenue protection as a function of order fill rate and SLA compliance. This visibility enables warehouse managers to optimize for profitability rather than for operational metrics in isolation, and to communicate operational investment needs in the financial language that supply chain leadership requires [1][2].

The compounding profitability effect of coordinated four-pillar AI deployment is the framework's most significant contribution to warehouse economics. Improved demand forecasting accuracy reduces safety stock requirements, freeing working capital and reducing obsolescence risk. Autonomous fulfillment reduces labor cost per unit, improving operating margin on each order shipped. AI slotting optimization increases throughput capacity within existing footprint, supporting revenue growth without proportional capital investment in facility expansion. Real-time operational intelligence reduces waste from reactive management, SLA breach recovery, and inefficient labor deployment. These four levers operate simultaneously and reinforce each other: better demand data improves robot routing optimization, slotting improvements amplify robot throughput, and operational intelligence distributes the benefits of all three pillars evenly across operational time periods rather than concentrating them in narrow optimal conditions [3][4].

7. Implementation Pathway and Challenges

A phased AI deployment sequence that aligns capability introduction with organizational data maturity and investment capacity substantially reduces both technical and operational risk. Phase 1—Demand Intelligence Foundation (months 1–6)—prioritizes AI demand forecasting and inventory optimization as the first AI capability, because it delivers the highest return on investment relative to implementation complexity and requires the lowest integration footprint: historical sales data, basic ERP connectivity, and a cloud-based ML forecasting service. This phase builds the demand data infrastructure—data pipelines, feature engineering processes, model governance protocols—that serves as the foundation for all subsequent phases. Phase 1 also generates the forecast accuracy improvements and inventory cost reductions that fund subsequent investment, creating a self-financing AI adoption pathway accessible to SMEs with constrained capital budgets [1].

Phase 2—Operational Intelligence Layer (months 4–9, overlapping with Phase 1)—deploys real-time dashboards and SLA prediction systems, capitalizing on the data infrastructure established in Phase 1. This phase transforms management visibility and converts reactive incident response into proactive performance management without requiring the capital investment of physical automation. The combination of improved demand intelligence and real-time operational visibility enables managers to demonstrate measurable service quality and cost improvements that build organizational confidence and stakeholder support for the more significant capital investments required in subsequent phases [6].

Phase 3—Autonomous Fulfillment (months 9–18)—introduces AMR-based picking and order fulfillment systems, enabled by the operational data foundation and demand intelligence developed in earlier phases. AMR routing and wave planning algorithms perform significantly better in data-rich environments where accurate demand

forecasts inform wave planning and real-time operational dashboards enable dynamic fleet reallocation. Workforce change management is the critical success factor in this phase: warehouse staff require structured transition programs that develop the system supervision, exception handling, and AI feedback contribution competencies that AI-adjacent warehouse roles demand. Organizations that invest in participatory AMR deployment—involving pickers and supervisors in system configuration and testing—consistently achieve faster adoption and more effective ongoing performance improvement [4].

Phase 4—Dynamic Layout Intelligence (months 15–24)—introduces continuous AI slotting optimization and digital twin simulation, leveraging the full operational data environment established in prior phases to generate slotting recommendations calibrated to actual velocity distributions rather than historical snapshots. Data integration prerequisites for this phase—unified WMS and AMR data feeds, complete inventory position visibility, real-time transaction streams—are largely satisfied by the infrastructure built in earlier phases, reducing the implementation effort relative to deploying this capability in isolation. For SMEs, cloud-based AI-as-a-service deployment models reduce the capital requirements of all four phases substantially, enabling module-by-module AI adoption without the enterprise software licensing and on-premises infrastructure investment that has historically constrained smaller operators' access to warehouse intelligence capabilities [2].

8. Conclusion

This article has proposed and evaluated an integrated four-pillar AI framework for warehouse operational improvement, demonstrating that demand forecasting intelligence, autonomous robotic fulfillment, AI-powered layout optimization, and real-time operational intelligence are not independent technology investments but synergistic capability pillars whose combined deployment generates compounding profitability improvements that exceed the sum of their individual contributions. The quantitative evidence base reviewed across all four pillars is consistent and substantial: forecast error reductions of 20–40%, picking throughput improvements of 3–5×, travel time reductions of 20–35%, and fulfillment cycle time reductions of 25–35% are achievable outcomes in well-executed deployments, each translating into direct financial value through reduced labor cost, lower inventory carrying cost, improved margin, or enhanced revenue through service quality differentiation.

The implications for warehouse operators are strategically significant. Warehouse management is undergoing a structural transition from a cost-containment function to a strategic competitive differentiator—and the organizations that close the AI capability gap earliest will establish fulfillment economics and service level advantages that become increasingly difficult for laggards to overcome as customer expectations continue to rise. The phased implementation pathway presented in this article makes this transition manageable across organizational scales, enabling SMEs to capture meaningful AI-driven efficiency gains through sequenced, self-financing investments that build toward full four-pillar deployment over twenty-four months. The framework positions AI not as a wholesale disruption requiring large upfront commitment, but as a progressive operational transformation that generates measurable return at each phase while building toward higher-order capability integration.

Future research should advance the framework in three directions. First, empirical longitudinal studies of four-pillar AI deployment outcomes across multi-year operational cycles would validate the compounding profitability model and identify the organizational and technical factors that determine whether performance gains are sustained as systems mature and operating conditions evolve. Second, multi-echelon AI optimization—extending the warehouse intelligence framework to encompass supplier networks, transportation lanes, and customer demand sensing in an integrated supply chain intelligence architecture—represents the next-order capability frontier with significant academic and practical value. Third, the application of generative AI to warehouse configuration design—using large language models and generative algorithms to synthesize optimal warehouse layouts, slotting configurations, and fulfillment workflows from operational parameters and constraint specifications—is an emerging capability direction that warrants structured research investment to establish its practical utility and implementation boundaries.

REFERENCES

1. W. A. Walter, K. Ahsan, and S. Rahman, "Application of artificial intelligence in demand planning for supply chains: a systematic literature review," *International Journal of Logistics Management*, vol. 36, no. 3, pp. 672–719, 2025. [Online]. Available: [https://www.emerald.com/ijlm/article/36/3/672/1246964/ Application-of-artificial-intelligence-in-demand](https://www.emerald.com/ijlm/article/36/3/672/1246964/Application-of-artificial-intelligence-in-demand)
2. R. F. de Assis et al., "Machine learning in warehouse management: a survey," *Procedia Computer Science*, vol. 232, pp. 2790–2799, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1877050924002734>
3. L. Zhen et al., "Optimizing warehouse operations with autonomous mobile robots," *Transportation Science*, vol. 59, no. 5, pp. 1130–1152, 2025. [Online]. Available: <https://pubsonline.informs.org/doi/abs/10.1287/trsc.2024.0800>
4. F. F. Rad et al., "Adoption of AI-based order picking in warehouse: benefits, challenges, and critical success factors," *Review of Managerial Science*, vol. 19, no. 11, pp. 3495–3540, 2025. [Online]. Available: <https://link.springer.com/article/10.1007/s11846-025-00858-1>
5. S. Mahmoudinazlou et al., "Deep reinforcement learning for dynamic order picking in warehouse operations," *Computers and Operations Research*, vol. 182, p. 107143, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0305054825001406>
6. "Intelligent warehousing: a machine learning and IoT framework for precision inventory optimization," *IEEE Access*, 2025. DOI: 10.1109/ACCESS.2025.3614679. [Online]. Available: <https://ieeexplore.ieee.org/document/11181053/>
7. "Storage location assignment for improving human-robot collaborative order-picking efficiency in robotic mobile fulfillment systems," *Sustainability*, vol. 16, no. 5, p. 1742, 2024. [Online]. Available: <https://www.mdpi.com/2071-1050/16/5/1742>