

A Framework for Vision-Based Obstruction Detection in High-Density Robotic Warehouse Systems: Design Principles and Deployment Outcomes

Livia Silveira Bezerra de Menezes

Engineering Project Manager

Efficient Plant Inc. (EPI)

Laval, Quebec, Canada

liviasbm@hotmail.com

Rogério Teixeira Mâsih

Adjunct Professor II in the Department of Production Engineering

Universidade Federal do Ceará

Fortaleza, Ceará, Brazil

rogeriom@ufc.br

Abstract

This paper proposes a framework for designing and deploying vision-based obstruction detection systems in high-density robotic warehouse environments. Traditional proximity-based detection approaches fail as product profiles diversify and operational envelopes expand — a challenge increasingly common as automated distribution platforms extend to new client segments. The proposed framework addresses this gap through three components: a cell-level detection architecture that replaces discrete sensor triggers with continuous visual pattern recognition; a product-profile-driven calibration methodology that tailors detection logic to the specific characteristics of each deployment context; and a release management integration layer that converts detection capability into active incident prevention. The framework is validated through a production deployment achieving a 99% detection rate, 98% incident prevention rate, and 75% reduction in network-wide push-through incidents across 26 distribution center sites. Five design principles are extracted from the deployment experience, providing generalizable guidance for practitioners designing detection systems for expanding automated warehouse platforms.

Keywords

vision-based detection, warehouse robotics, obstruction detection, computer vision, industrial engineering.

1. Introduction

Automated warehouse robotic systems are subject to a class of failure mode — push-through incidents — that traditional sensor-based detection approaches handle inadequately as platforms scale and diversify. A push-through incident occurs when a robotic lifting mechanism continues to operate despite an obstruction in the lift cell, resulting in product damage, equipment damage, or operational disruption.

The inadequacy of conventional approaches becomes acute during product profile expansion: when an automated platform is extended to serve clients with product characteristics outside the original design envelope, existing sensor configurations are no longer calibrated to the actual operational conditions. This is not a marginal edge case — it is a recurring engineering challenge for any platform seeking to grow its addressable market.

This paper proposes a structured framework for designing vision-based obstruction detection systems that addresses this challenge. The framework's primary contribution is not the specific system deployed, but the generalizable framework derived from that deployment — a set of architectural decisions, calibration principles, and integration patterns that practitioners can apply to their own detection design challenges.

1.1 Objectives

This paper pursues three objectives: (1) to propose a structured, three-component framework for vision-based obstruction detection that generalizes beyond a single sensor implementation; (2) to validate that framework through a staged production deployment, measuring detection rate, prevention rate, and network-wide incident reduction; and (3) to extract generalizable design principles from the deployment experience that practitioners can apply when designing detection systems for their own expanding automated warehouse platforms. These objectives are revisited explicitly in the Conclusion.

2. Literature Review

The relevant literature spans three areas: the limitations of proximity-based detection, the development of vision-based approaches in industrial automation, and the integration of detection outputs into operational decision systems. This section reviews each before identifying the gap this paper addresses.

Proximity-based detection in robotic warehouse systems relies on discrete sensor triggers calibrated for a specific product profile and operational envelope (Siciliano et al. 2009). This approach has well-documented limitations in high-variability environments: fixed-position sensors cannot adapt to changes in product geometry, load distribution, or stacking configuration without physical reconfiguration and recalibration (Monkman et al. 2007). The limitation becomes particularly acute in multi-client platform deployments, where a single robotic system must handle diverse product categories across multiple distribution environments; sensor configurations optimized for one client's product mix may fail to detect obstruction conditions arising from another client's products (de Koster et al. 2007).

Computer vision applications in industrial automation have expanded significantly with advances in edge computing hardware and deep learning techniques (LeCun et al. 2015; Krizhevsky et al. 2012). Vision-based systems offer a fundamentally different detection paradigm: continuous evaluation of the visual state of the operating environment through pattern recognition, rather than discrete threshold triggering. Applications of vision-based approaches in warehouse environments have been reported for inventory management (Fragapane et al. 2021), pick accuracy verification (Wurman et al. 2008), and safety monitoring (Boysen et al. 2019). However, the specific application of vision-based detection to lift obstruction prevention in high-density robotic environments — and the design framework for such systems — is not well documented in the existing literature.

A detection system that identifies an obstruction condition but cannot act on that information delivers limited operational value. The integration of detection outputs into operational release management frameworks — enabling real-time cell exclusion based on detected conditions — is a critical architectural dimension not addressed in prior vision-based detection literature (Humble and Farley 2010). This integration layer, together with the broader warehouse-automation operating context documented by Roodbergen and Vis (2009) and Tompkins et al. (2010), motivates the third component of the framework proposed in this paper.

Taken together, existing literature establishes the limitations of proximity-based detection, the technical foundations of vision-based pattern recognition, and the general principle that detection value depends on downstream integration — but no prior work combines cell-level vision architecture, product-profile-specific calibration, and formal release-management integration into a single validated framework for lift-obstruction prevention. This paper addresses that gap.

3. Methods

The proposed framework consists of three components, illustrated in Figure 1: a cell-level detection architecture, a product-profile-driven calibration methodology, and a release management integration layer.




Figure 1. Proposed vision-based obstruction detection framework (three-component architecture).

3.1 Component 1: Cell-Level Detection Architecture

The framework prescribes detection at the individual lift cell level rather than at the system or zone level. Cell-level detection provides three advantages: it localizes detection to the specific cell where an obstruction condition exists, enabling targeted response without disrupting adjacent operations; it allows detection logic to be calibrated for the specific geometry and product profile of each cell configuration; and it integrates naturally with cell-level release management. Camera placement is determined by the operational envelope of the installed end-of-arm tooling (EOAT) configuration; as EOAT is adapted for different product profiles, camera placement must be re-evaluated to ensure coverage of the full operational envelope.

3.2 Component 2: Product-Profile-Driven Calibration

The framework prescribes a calibration methodology in which detection logic is developed and validated against the specific product and load profiles of the target deployment context, not against a generic baseline profile. This methodology involves three steps: profile characterization (systematic documentation of product categories, weight ranges, geometric variability, and stacking configurations); logic training (development of detection logic against representative data, with attention to edge cases at the limits of weight and geometry); and threshold validation (validation against both the target profile and adjacent profiles that may be introduced as the platform expands). This methodology treats detection logic as a product-profile artifact that must be updated when the product profile changes, rather than a one-time configuration.

3.3 Component 3: Release Management Integration

The framework's third component prescribes formal integration between detection outputs and the operational release management system: when the detection system identifies a cell-level obstruction condition, that output must carry formal authority to exclude the affected cell from active operation pending resolution. This integration transforms detection from a monitoring function into a prevention function. For push-through incident prevention specifically,

where the consequence of the incident is product or equipment damage, prevention rather than documentation is the operational objective.

4. Data Collection

The framework was instantiated and validated in the context of an automated warehouse robotic platform serving North American retail distribution networks, operating across hundreds of lift cells per distribution center and processing high volumes of product movements per shift across a diverse product mix, including conventional packaged goods and, following a platform expansion, heavy beverage distribution products. The product profile expansion was the precipitating event for the detection framework's development, since existing proximity-based detection thresholds were not adequate for the expanded product mix.

Data collection followed a two-stage protocol. Stage 1 — single-site pilot: full deployment of all three framework components at a single distribution center site under production conditions, with measurement of detection rate and prevention rate against pre-defined KPI targets. Stage 2 — network rollout: staged deployment across 26 distribution center sites, with monitoring gates between deployment batches and measurement of network-wide incident rates. The staged rollout approach used the gate-based release management framework described in related work (Menezes 2026), which provided the change control structure for network-scale deployment.

5. Results and Discussion

5.1 Numerical Results

The 99% detection rate confirms that the cell-level camera architecture provides adequate coverage of the operational envelope across the full product mix, including the expanded beverage distribution product category. The 98% prevention rate confirms that the release management integration layer is effective at converting detection events into operational interventions before incidents complete (Table 1).

Table 1. Key performance outcomes from the production deployment of the proposed detection framework

Metric	Result	Context
Detection rate (pilot site)	99%	Production conditions, full product mix
Incident prevention rate	98%	Detected conditions resolved before damage
Network incident reduction	75%	vs. pre-framework baseline, 26 sites
Sites deployed	26	North American distribution centers

5.2 Graphical Results

Figure 2 presents the three headline outcome metrics together. The 75% reduction in network-wide push-through incidents represents the combined effect of the detection framework and the concurrent release management framework — neither alone would have achieved this outcome. The interaction between detection capability and release management authority is a key finding with design implications, discussed further in Section 5.3.




Figure 2. Detection rate, incident prevention rate, and network-wide incident reduction from the production deployment.

5.3 Proposed Improvements

Five design principles are extracted from the deployment experience as generalizable guidance for practitioners designing detection systems for expanding automated warehouse platforms.

- **Principle 1 — Detection architecture must be product-profile-specific, not generic.** Sensor configurations and detection logic designed for one product profile will degrade as the operational envelope expands; this is an architectural issue, not a calibration issue.
- **Principle 2 — Camera placement is an EOAT-dependent decision.** The camera field of view must cover the full operational envelope of the installed EOAT configuration; these are not independent engineering decisions.
- **Principle 3 — Cell-level detection enables targeted response without operational disruption.** Cell-level detection localizes the response to the affected unit, preserving throughput at adjacent cells.
- **Principle 4 — Release management integration is the difference between detection and prevention.** A detection system without formal authority to act on detection outputs is a monitoring system, not a prevention system.
- **Principle 5 — Staged validation is required before network-scale deployment.** Single-site validation does not cover the configuration variability introduced at network scale; a staged rollout — pilot, limited expansion, network release — is necessary to identify configuration-specific issues before they propagate.

5.4 Validation

The framework was validated through the two-stage protocol described in Section 4: a single-site pilot under full production conditions, followed by a staged rollout across 26 distribution center sites with monitoring gates between batches. This paper presents findings from a single deployment context; the 75% network-wide reduction is the joint outcome of the detection framework and the concurrent release management framework, and the relative contribution of each cannot be fully separated from this dataset alone. Positioned against prior work, the achieved detection and prevention rates substantially exceed what proximity-based approaches typically sustain under expanding product-profile conditions (Siciliano et al. 2009; Monkman et al. 2007), and the finding that integration with formal release authority is necessary for prevention, not merely detection, extends the continuous-delivery literature's emphasis on automated release gating (Humble and Farley 2010) into a physical-system, safety-relevant context that literature does not address.

6. Conclusion

This paper proposed a three-component framework for vision-based obstruction detection in high-density robotic warehouse environments, validated through a production deployment achieving a 99% detection rate, 98% incident prevention rate, and 75% reduction in network-wide push-through incidents across 26 distribution center sites. Returning to the three objectives stated in Section 1.1: the three-component framework was presented in full

(Objective 1); it was validated through a two-stage production deployment against pre-defined KPI targets (Objective 2); and five generalizable design principles were extracted from the deployment experience (Objective 3).

The framework's primary contribution is the identification of the architectural and methodological decisions that determine detection system effectiveness in expanding platform contexts. For practitioners managing automated warehouse platform expansions, the key implication is that detection architecture must be treated as a product-profile-specific engineering challenge, not a one-time configuration.

References

- Boysen, N., de Koster, R. and Weidinger, F., Warehousing in the e-commerce era: a survey, *European Journal of Operational Research*, vol. 277, no. 2, pp. 396-411, 2019.
- de Koster, R., Le-Duc, T. and Roodbergen, K. J., Design and control of warehouse order picking: a literature review, *European Journal of Operational Research*, vol. 182, no. 2, pp. 481-501, 2007.
- Fragapane, G., de Koster, R., Sgarbossa, F. and Strandhagen, J. O., Planning and control of autonomous mobile robots for intralogistics: literature review and research agenda, *European Journal of Operational Research*, vol. 294, no. 2, pp. 405-426, 2021.
- Humble, J. and Farley, D., *Continuous Delivery: Reliable Software Releases through Build, Test, and Deployment Automation*, Addison-Wesley, Boston, 2010.
- Krizhevsky, A., Sutskever, I. and Hinton, G., ImageNet classification with deep convolutional neural networks, *Advances in Neural Information Processing Systems*, vol. 25, pp. 1097-1105, 2012.
- LeCun, Y., Bengio, Y. and Hinton, G., Deep learning, *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.
- Menezes, L. S. B., A gate-based release management framework for multi-site robotic system deployment, *Proceedings of the 7th Latin American Conference on Industrial Engineering and Operations Management*, Panama City, Panama, 2026.
- Monkman, G. J., Hesse, S., Steinmann, R. and Schunk, H., *Robot Grippers*, Wiley-VCH, Weinheim, 2007.
- Muller, M., *Essentials of Inventory Management*, HarperCollins Leadership, Nashville, 2019.
- Richards, G., *Warehouse Management: A Complete Guide to Improving Efficiency and Minimizing Costs in the Modern Warehouse*, Kogan Page, London, 2017.
- Roodbergen, K. J. and Vis, I. F. A., A survey of literature on automated storage and retrieval systems, *European Journal of Operational Research*, vol. 194, no. 2, pp. 343-362, 2009.
- Siciliano, B., Sciavicco, L., Villani, L. and Oriolo, G., *Robotics: Modelling, Planning and Control*, Springer, London, 2009.
- Tompkins, J. A., White, J. A., Bozer, Y. A. and Tanchoco, J. M. A., *Facilities Planning*, 4th ed., Wiley, Hoboken, 2010.
- Wurman, P. R., D'Andrea, R. and Mountz, M., Coordinating hundreds of cooperative, autonomous vehicles in warehouses, *AI Magazine*, vol. 29, no. 1, pp. 9-19, 2008.
- Zhang, Q., Vonderembse, M. A. and Lim, J. S., Manufacturing flexibility: defining and analyzing relationships among competence, capability, and customer satisfaction, *Journal of Operations Management*, vol. 21, no. 2, pp. 173-191, 2003.

Biographies

Livia Silveira Bezerra de Menezes is an Engineering Project Manager specializing in R&D program management for supply chain automation and warehouse robotics systems. She holds a Master of Industrial Engineering from Concordia University (Montreal) and a BSc in Mechanical Production Engineering from Universidade Federal do Ceara (Brazil). She is a PMP-certified Project Management Professional, a Lean Six Sigma Green Belt (McGill University), and a licensed engineer (Ordre des Ingenieurs du Quebec — OIQ).

Dr Rogério Teixeira Mâsih is an Adjunct Professor II in the Department of Production Engineering, Center of Technology, at the Universidade Federal do Ceará (UFC), Brazil. He holds a BSc in Mechanical Engineering from UFC (1997) and an MSc (1999) and PhD (2005) in Production Engineering from the Universidade Federal de Santa Catarina (UFSC), where his doctoral research focused on strategic competency mapping and its integration with the Balanced Scorecard for competency-based human resource management. He has taught and advised graduate and undergraduate students in production engineering, logistics, and operations management at UFC since 2006, and served as Associate Dean for University Extension (Pró-Reitor Adjunto de Extensão) from 2013 to 2023 and as

*Proceedings of the 7th Latin American Conference on Industrial Engineering and Operations Management
Panama City, Panama, August 4-6, 2026*

Coordinator of UFC's Internship Agency from 2012 to 2023. His research and teaching interests include strategic cost management, performance measurement systems, logistics, and production planning and control.