

VisionTwin-RL: Digital-Twin-Guided Reinforcement Learning for Adaptive Visual Inspection Systems

Ahmad Taufik Al Afkari ¹

Department of Computer Science, Universitas Islam Negeri Sumatera Utara, Indonesia, Medan
Corresponding author: Ahmad Taufik Al Afkari, ahmadtaufikalafkari@uinsu.ac.id

Article Info

Article history:
Received: May 25, 2026
Revised: July 1, 2026
Accepted: July 11, 2026

Keywords:

Digital Twin
Reinforcement Learning
Visual Inspection
Smart Manufacturing
Quality Control
Adaptive Acquisition

ABSTRACT

Industrial visual inspection must balance defect-detection quality, cycle time, and inspection cost under changing product and imaging conditions. Digital twins provide a risk-free environment for testing control policies, while reinforcement learning can adapt decisions to observed state. This paper proposes VisionTwin-RL, a digital-twin-guided inspection policy that selects among fast, balanced, and high-fidelity acquisition modes according to estimated defect risk and image degradation. A controlled nine-state digital twin was constructed from three defect-risk levels and three blur levels. A tabular reinforcement-learning agent was trained for 5,500 episodes using a reward that combines detection probability, inspection cost, and the consequence of a missed defect. The learned policy was compared with two static inspection strategies. VisionTwin-RL achieved a higher detection–cost balance and the highest mean reward in the simulated evaluation. The results demonstrate the computational value of state-dependent acquisition rather than a fixed inspection setting. The study is a proof-of-concept digital-twin experiment; the detection probabilities are controlled simulation parameters and must be replaced by measurements from a real vision model and physical inspection cell before industrial deployment claims are made.

This is an open access article under the [CC BY-SA](#) license.



Corresponding Author:

Ahmad Taufik Al Afkari
Department of Computer Science, Universitas Islam Negeri Sumatera Utara, Indonesia, Medan
Email: ahmadtaufikalafkari@uinsu.ac.id

1. INTRODUCTION

Smart manufacturing increasingly relies on automated inspection, digital twins, and learning-based control. Recent research spans digital-twin architectures, quality prediction, robot twins, virtual commissioning, visual defect detection, and reinforcement-learning-driven acquisition [1]–[11], [19]–[25]. These technologies create an opportunity to move inspection from a fixed camera-and-threshold pipeline toward an adaptive decision process.

Visual quality systems face a recurring trade-off. High-fidelity acquisition may improve detectability but increases cycle time, computation, energy consumption, or robot motion. Fast inspection reduces cost but can miss small or ambiguous defects. Recent defect-detection work addresses imbalance, domain shifts, weak supervision, micro-defects, and changing production backgrounds [12]–[18]. However, many pipelines still optimize the detector separately from the acquisition or inspection policy.

Digital twins can provide a safe environment in which inspection actions are evaluated before being applied to physical equipment. Reinforcement learning (RL) is particularly suitable when the best action depends on state and when direct trial-and-error on production hardware is expensive. Recent reviews and inspection studies show growing convergence between RL, digital twins, robotics, and adaptive acquisition [1], [3], [4], [22].

VisionTwin-RL operationalizes this convergence with a minimal digital twin containing defect-risk and image-quality state variables. The policy chooses one of three acquisition modes. Its reward favors successful detection while penalizing inspection cost and missed-defect risk. The approach is intentionally simple enough to expose the decision logic and can later be replaced by deep RL once state spaces include camera pose, illumination, product type, robot kinematics, and detector confidence.

The study evaluates whether a learned state-dependent policy can dominate static acquisition strategies in a controlled environment. The simulation does not claim to reproduce any particular factory; instead, it establishes a reproducible methodological scaffold for integrating measured vision-model performance into a digital twin.

2. METHOD

2.1. Digital-twin state and action space

The digital twin contains nine discrete states formed by three estimated defect-risk levels and three image-degradation levels. Actions correspond to fast, balanced, and high-fidelity inspection. Each action has a distinct acquisition cost and baseline detection probability. Blur decreases detection probability, whereas high-fidelity acquisition compensates more strongly when defect risk is high.

$$s_t = (risk_t, blur_t), \quad a_t \in \{fast, balanced, high-fidelity\}.$$

2.2. Reward design

The reward combines detection probability, acquisition cost, and an expected miss penalty that increases with risk. This prevents the agent from choosing the highest-fidelity action for every state and converts inspection into a quality–cost optimization problem.

$$r_t = p_{det}(s_t, a_t) - c(a_t) - (1 - p_{det}(s_t, a_t)) \cdot \lambda_{miss}(risk_t).$$

2.3. Reinforcement learning

A tabular Q-learning agent is trained for 5,500 episodes with epsilon-greedy exploration. The learning rate is 0.15, the discount factor is 0.10, and epsilon decays to a minimum of 0.03. Because the current state is small, tabular learning provides a transparent baseline; a deep policy is not required to demonstrate the adaptive acquisition principle.

2.4. Evaluation

The learned policy is compared with two fixed strategies: always-fast and always-high-fidelity. Four thousand randomized evaluation states are sampled with a fixed seed. Metrics are empirical detection rate, normalized inspection cost, and mean reward. The parameters describe a controlled twin rather than measurements from a real camera system.

3. RESULTS AND DISCUSSION

Table 1 reports the quality–cost trade-off. The always-high policy increases detection but incurs the largest cost. The always-fast policy is inexpensive but loses detection performance. VisionTwin-RL selects modes conditionally and obtains the highest mean reward in the controlled twin.

The result is important because it separates detector design from inspection policy design. A more accurate detector does not automatically imply an efficient inspection system. When the acquisition mode can be changed, the controller should allocate expensive sensing or robot motion only to states where expected quality gain justifies the cost.

Policy	Detection rate	Normalized cost	Mean reward
Static-fast	0.617	0.040	0.179
Static-high	0.916	0.350	0.490
VisionTwin-RL	0.881	0.249	0.521

Table 1. Detection–cost performance of static and adaptive policies.




Figure 1. Detection rate and normalized inspection cost in the controlled digital twin.

Figure 1 visualizes detection rate and normalized cost. The proposed method occupies the middle-cost region while retaining strong detection. This is the intended role of the digital twin: it allows a policy to explore trade-offs without creating defective products, slowing a physical line, or risking unsafe robot behavior.

A practical implementation would replace the hand-specified detection function with a calibrated vision model evaluated across camera pose, illumination, blur, surface texture, and product families. The twin would also include action duration and robot-reachability constraints. The reinforcement-learning formulation remains unchanged, but state estimation and reward calibration become empirical engineering tasks.

3.1. Threats to validity and limitations

Detection probabilities, blur effects, defect risk, and inspection costs are synthetic parameters. The study therefore evaluates the control architecture, not industrial accuracy. Tabular Q-learning assumes a compact Markov state and does not represent complex visual observations directly. The digital twin has no physical synchronization error, network latency, robot collision model, or sim-to-real gap. A journal submission claiming industrial validation should couple VisionTwin-RL to a real defect detector, report physical cycle times and energy/cost measures, calibrate the twin against observed data, and evaluate policy safety and transferability across products.

4. CONCLUSION

VisionTwin-RL demonstrates a digital-twin-guided reinforcement-learning approach for state-dependent visual inspection. The controlled study shows that a learned policy can allocate high-fidelity inspection selectively, preserving detection quality while avoiding the cost of always using the most expensive mode. The contribution is an integrated decision formulation rather than a new defect detector. Future research should connect the twin to measured detector calibration curves, incorporate continuous camera and robot actions, model temporal production dynamics, enforce safety constraints, and validate sim-to-real transfer on a physical inspection cell.

REFERENCES

- [1] M. P. Ramos and A. Busboom, "Integration of digital twins with reinforcement learning in industry: A systematic review," in *Proc. IEEE ETFA*, 2025, doi: 10.1109/ETFA65518.2025.11205657.
- [2] P. Nguyen, M. Kim, E. Nichols, and H.-S. Yoon, "AI-driven digital twins for manufacturing: A review across hierarchical manufacturing system levels," *Sensors*, vol. 26, no. 1, art. 124, 2026, doi: 10.3390/s26010124.
- [3] Y. M. Alginahi, O. Sabri, and W. Said, "Reinforcement learning for industrial automation: A comprehensive review of adaptive control and decision-making in smart factories," *Machines*, vol. 13, no. 12, art. 1140, 2025, doi: 10.3390/machines13121140.
- [4] S. Chen et al., "Enhancing digital twins with deep reinforcement learning: A use case in maintenance prioritization," in *Proc. Winter Simulation Conf.*, 2024, doi: 10.1109/WSC63780.2024.10838867.
- [5] H. A. Faqeer and S. H. Khajavi, "Digital twin and computer vision combination for manufacturing and operations: A systematic literature review," *Applied Sciences*, vol. 15, no. 18, art. 10157, 2025, doi: 10.3390/app151810157.
- [6] D. Y. Patrashko and V. Gurau, "Machine learning-powered vision for robotic inspection in manufacturing: A review," *Sensors*, vol. 26, no. 3, art. 788, 2026, doi: 10.3390/s26030788.
- [7] P. Link et al., "Real-time quality prediction and local adjustment of friction with digital twin in sheet metal forming," *Robotics and Computer-Integrated Manufacturing*, vol. 91, art. 102848, 2025, doi: 10.1016/j.rcim.2024.102848.
- [8] X. Fu, S. Li, H. Song, and Y. Lu, "Digital twin-driven multi-scale characterization of machining quality: Current status, challenges, and future perspectives," *Robotics and Computer-Integrated Manufacturing*, vol. 93, art. 102902, 2025, doi: 10.1016/j.rcim.2024.102902.
- [9] M. H. Zafar, E. F. Langás, and F. Sanfilippo, "Exploring the synergies between collaborative robotics, digital twins, augmentation, and Industry 5.0 for smart manufacturing: A state-of-the-art review," *Robotics and Computer-Integrated Manufacturing*, vol. 89, art. 102769, 2024, doi: 10.1016/j.rcim.2024.102769.
- [10] S. Wang et al., "A deep learning-enhanced digital twin framework for improving safety and reliability in human-robot collaborative manufacturing," *Robotics and Computer-Integrated Manufacturing*, vol. 85, art. 102608, 2024, doi: 10.1016/j.rcim.2023.102608.
- [11] H. Ni et al., "Digital twin-driven virtual commissioning for robotic machining enhanced by machine learning," *Robotics and Computer-Integrated Manufacturing*, vol. 93, art. 102908, 2025, doi: 10.1016/j.rcim.2024.102908.
- [12] X. Liang et al., "Adaptive shape imitation and selective semantic guidance for industrial surface defect detection," *Expert Systems with Applications*, vol. 279, art. 127334, 2025, doi: 10.1016/j.eswa.2025.127334.
- [13] C. Wang, X. Wei, and X. Jiang, "An automated defect detection method for optimizing industrial quality inspection," *Engineering Applications of Artificial Intelligence*, vol. 127, art. 107387, 2024, doi: 10.1016/j.engappai.2023.107387.
- [14] D. Bai et al., "Surface defect detection methods for industrial products with imbalanced samples: A review of progress in the 2020s," *Engineering Applications of Artificial Intelligence*, vol. 130, art. 107697, 2024, doi: 10.1016/j.engappai.2023.107697.
- [15] "Robust weakly supervised product surface defect segmentation based on guided cropping and inpainting extension," *Applied Soft Computing*, vol. 184, art. 113661, 2025, doi: 10.1016/j.asoc.2025.113661.
- [16] Q. Qin et al., "Robot digital twin systems in manufacturing: Technologies, applications, trends and challenges," *Robotics and Computer-Integrated Manufacturing*, vol. 97, art. 103103, 2026, doi: 10.1016/j.rcim.2025.103103.
- [17] "Micro-defect VariFocal Network: Channel attention and spatial feature fusion for turbine blade surface micro-defect detection," *Engineering Applications of Artificial Intelligence*, vol. 133, art. 108075, 2024, doi: 10.1016/j.engappai.2024.108075.
- [18] "Test-time feature caching network for cross-domain multilayer ceramic capacitors defect detection," *Applied Soft Computing*, vol. 179, art. 113380, 2025, doi: 10.1016/j.asoc.2025.113380.
- [19] C. Latsou et al., "A unified framework for digital twin development in manufacturing," *Advanced Engineering Informatics*, vol. 62, art. 102567, 2024, doi: 10.1016/j.aei.2024.102567.
- [20] C. Zhang et al., "Hybrid mechanism and data-driven digital twin model for assembly quality traceability and optimization of complex products," *Advanced Engineering Informatics*, vol. 62, art. 102707, 2024, doi: 10.1016/j.aei.2024.102707.
- [21] A. Gautam et al., "IIoT-enabled digital twin for legacy and smart factory machines with LLM integration," *Journal of Manufacturing Systems*, vol. 80, pp. 511–523, 2025, doi: 10.1016/j.jmsy.2025.03.022.

- [22] J.-P. Kaiser et al., “Adaptive acquisition planning for visual inspection in remanufacturing using reinforcement learning,” *Journal of Intelligent Manufacturing*, vol. 36, pp. 4867–4893, 2025, doi: 10.1007/s10845-024-02478-0.
- [23] X. Fu, H. Song, S. Li, and Y. Lu, “Digital twin technology in modern machining: A comprehensive review of research on machining errors,” *Journal of Manufacturing Systems*, vol. 79, pp. 134–161, 2025, doi: 10.1016/j.jmsy.2025.01.005.
- [24] B. Kang, C. Park, H. Kim, and S. Hong, “A digital twin calibration for an automated material handling system in a semiconductor fab,” *Journal of Manufacturing Systems*, vol. 80, pp. 1013–1028, 2025, doi: 10.1016/j.jmsy.2025.04.015.
- [25] X. Tong, J. Bao, and F. Tao, “Co-evolutionary digital twins: A multidimensional dynamic approach to digital engineering,” *Advanced Engineering Informatics*, vol. 61, art. 102554, 2024, doi: 10.1016/j.aei.2024.102554.