UNIVERSITY OF WEST ATTICA

eISSN:3136-7658

Applied Manufacturing Intelligence Journal

Vol. 5, No. 1, January, 2026



Inventory State Reconstruction in Smart Warehouses Using a GraphGPS-VAE Model with Cross-Domain Fusion

Kristo Pärn[1,*], Gert Berg[1]

[1] *Department of Artificial Intelligence, Narva College of Economics, Narva, 20304, Estonia*

Received 18 October 2025; Revised 01 January 2026; Accepted 04 January 2026; Published 09 January 2026

Abstract

This study presents a GraphGPS-VAE model with cross-domain fusion for inventory-state reconstruction in smart warehouses. The task is to infer a reliable shelf-level state from RFID events, visible shelf images, weight-sensor records, warehouse-management logs, and picking records when these sources are incomplete or inconsistent. Shelf locations, stock-keeping units, inbound batches, outbound operations, and sensing devices are first organized as a heterogeneous shelf-item-sensor graph. The GraphGPS encoder combines local shelf adjacency with aisle-level and zone-level dependencies, while the VAE branch reconstructs expected quantity, location consistency, batch status, and sensor agreement. A cross-domain fusion module aligns transferable stock-flow patterns while preserving target-warehouse layout, turnover rhythm, and sensing residuals. The model is evaluated on a simulated dataset containing 14 warehouses, 132 storage zones, 48,000 shelf-state windows, 1.92 million sensing events, and six inventory-state labels. Compared with rule-based checking, isolated visual recognition, graph convolutional reconstruction, and VAE without domain fusion, the proposed method reduces quantity mean absolute error to 2.18 units, improves location-consistency accuracy to 92.7%, increases abnormal-state recall to 0.884, and reduces manual review load by 21.6%. These results suggest that graph-based cross-domain reconstruction can support exception tracing, replenishment review, and inventory auditing. The method is therefore positioned as decision support for warehouse staff rather than an automatic stock-account correction tool.

Keywords: Smart Warehouse; Inventory State Reconstruction; GraphGPS-VAE; Cross-Domain Fusion; Anomaly Detection

1. Introduction

Smart warehouse inventory management increasingly depends on integrated sensing, because modern facilities must track item quantity, shelf location, batch identity, picking status, and replenishment time within one operational record [1]. Inventory errors may delay orders, create overstock, reduce customer-service reliability, and distort financial stock records [2]. Multi-source perception can reduce blind spots, but RFID evidence, visual inspection, weight sensing, and WMS logs often disagree under dense storage, temporary holding, or delayed posting conditions [3]. Recent digital-warehouse research has also shown that site-to-site domain differences weaken models trained in one layout when they are deployed in another warehouse [4]. Inventory reconstruction should therefore be treated as a reliability problem that connects sensing evidence, operational history, and warehouse-specific context.

The inventory-state reconstruction problem is more difficult than ordinary item classification. Vision-based inventory research [5] shows that a camera may confirm that a carton is present, yet it may fail to verify the correct batch when labels are blocked by stacked goods. An RFID reader may detect a tag, but in a narrow aisle the tag can be associated with the wrong neighboring bin. A weight sensor may indicate a quantity change while remaining unable to distinguish replenishment, picking, or temporary relocation. WMS logs provide process records, but manual scan

* Corresponding author: kristo.p@nek.ee

https://doi.org/10.68406/amij.2026.vol5iss1np2:16-29

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

omissions and delayed posting can make the digital account lag behind the physical shelf. These examples show why a single data source cannot be treated as final inventory truth.

Graph-based models are suitable for the relational features in inventory data. Shelves are connected to aisles, bins, items, batches, operators, sensors, inbound orders, outbound orders and exception tickets. Cross-domain fusion is also needed because of different warehouse shelf spacings, product turnover rates, scanner densities, lighting, robot path designs and manual operation routines. A model that only memorizes source-domain patterns may perform well in a familiar warehouse but fail in a new cold-storage zone, a high-turnover e-commerce site, or a mixed manual-automation warehouse. Therefore, the present study treats smart warehouse inventory state as a reconstructable graph signal rather than a flat table or a single image label.

This paper is organised as follows. Section 2 reviews smart warehouse inventory perception, cross-domain representation learning and graph-based variational reconstruction models. Section 3 introduces the proposed GraphGPS-VAE model, which includes multi-source inventory graph construction, cross-domain fusion, GraphGPS encoding, variational reconstruction and review-level output. Section 4 presents the simulated dataset, baseline comparison, transfer evaluation, ablation analysis and field-oriented warning results. Section 5 is the main results and discussions, along with deficiencies and future research directions.

2. Related Work

2.1. Smart Warehouse Inventory Perception

Smart warehouse inventory perception combines operating data and sensing equipment to describe the physical and procedural condition of stored items. Research on data warehouses and machine-learning platforms [6] shows that sensing streams must be organized with process records before they can support trustworthy operational judgement. Vision systems identify occupancy, product type, and clear placement faults, but they are affected by occlusion, glare, stacked cartons, and uneven label orientation. RFID and IoT devices support object tracking, but reader overlap, liquid goods, and dense metal shelves can produce ambiguous tag-location evidence. WMS records document inbound, outbound, transfer, and cycle-count actions, while weight and conveyor sensors provide quantity-change evidence. Because each source fails in a different way, direct feature concatenation may amplify noise rather than improve reconstruction. A useful inventory model must learn which evidence source is reliable under each operating condition.

An inventory audit is fundamentally a reconstruction task rather than a simple detection task. Conventional checks confirm whether a stock-keeping unit exists, whether its quantity matches the account, and whether the location code is correct. Smart warehouses must additionally resolve batch mismatches, temporary storage, duplicate scans, outdated inbound confirmations, and unposted picking movements. Supply-chain visibility research [7] shows that inventory data should be connected to replenishment, allocation, and exception-management decisions. Digital twin studies [8] also indicate that the observed operational state should be compared with expected system behavior. These findings support a model that reconstructs a consistent shelf-level state from multiple weak signals.

2.2. Cross-Domain Representation Learning

The issue of training-deployment environment mismatch is resolved by cross-domain representation learning. In a smart warehouse, the target domain has different shelf configurations, a different percentage of fast-moving items, and fewer cycle-count confirmations than the source domain, which could include complete labels from an older facility. Feature alignment is acceptable when the fundamental goals are the same but the observation distributions are different, according to domain adaptation research [9]. This paper's common objective is inventory-state reconstruction, with aisle layout, scanner coverage, storage density, illumination, product mix, and operating schedule being domain-specific elements.

Domain shift in warehouses is not a single statistical change. It can be caused by mobile robots, temporary workforce changes, new packaging materials, altered picking routes, cold-chain handling, or large order surges. Research on privacy-preserving operational learning [10] also shows that warehouse data may need to be learned across organizations or sites where raw-data sharing is restricted. Although this study does not build a federated system, the motivation is similar: transferable stock-flow patterns should be learned without forcing every warehouse to look identical. The cross-domain fusion module therefore retains both useful source-domain evidence and target-domain residual information. This distinction matters because a new fast-picking zone should not be pushed back toward a slower source-domain pattern.

2.3. Graph-Based Variational Reconstruction Models

Graph neural networks can explicitly describe item-shelf-sensor interactions, making them suitable for inventory data. Recent graph-learning research [11] shows that local message passing preserves neighborhood structure, while

positional and global attention mechanisms capture longer-range interdependence. In warehouses, local interactions include nearby bins, shared cameras, and adjacent RFID coverage; global links include aisle routes, zone-level replenishment policies, and cross-zone picking sequences. GraphGPS-style modeling [12] is therefore attractive because it combines local propagation with global representation, allowing the model to learn both shelf adjacency and warehouse-wide operational structure.

When it is easier to observe the normal condition than all of the aberrant situations, variational reconstruction models are typically used. In inventory management, routine picking processes, accurate item-location matching, and typical shelf statuses are more common than exceptional circumstances. Uncertainty-aware learning [13] can differentiate between weak or absent evidence and stable anomalous states, and VAE-based anomaly detection research [14] supports the use of reconstruction residuals as warning indications. For the present issue, the VAE module reconstructs the predicted item number, position consistency, batch relationship, and sensor agreement.

The alignment component is further motivated by cross-domain graph representation learning [15], as warehouses differ in layout, device density, and operation rhythm yet share stock-flow logic. Because inventory exceptions are frequently relational—a single shelf, an item tag, or an order event may seem normal in isolation but become suspect when related to other evidence—graph anomaly detection research [16] is also pertinent.

Studies on warehouse digitalisation [17] and multimodal fusion [18] demonstrate that integrated evidence, as opposed to individual sensor outputs, is necessary for the operational system. The research supports the development of a reconstruction model that may produce interpretable inventory review findings by combining shelf position, item identification, order procedure, and sensing uncertainty.

Reconstruction outputs should be connected to review actions rather than merely prediction scores, according to operational decision-support research [19]. The need for traceable warning evidence has been further underlined by research on explainable artificial intelligence [20]. As a result, the suggested model logs the responsible shelf, sensing source, operation window, and reconstruction residual.

3. Proposed Method

3.1. Multi-Source Inventory Graph Construction

Initially, a heterogeneous inventory graph is constructed from warehouse observations. Nodes include shelf bins, stock-keeping units, inbound batches, outbound orders, transfer jobs, sensing devices, warehouse zones, and exception tickets. Edges are formed when entities share physical location, process continuity, batch history, temporal co-occurrence, sensing coverage, or zone dependency. For example, a shelf node is linked to order nodes through picking or replenishment events, to batch nodes through inbound history, to sensor nodes through observation coverage, and to item nodes through current placement. The graph is not just a flat WMS export; it preserves relationships among physical evidence, process evidence, and sensing evidence. When sources conflict, the model keeps the disagreement as reconstruction evidence rather than immediately modifying the account record.

$$G = (V, E, X) \tag{1}$$

where $\mathbf{G}$ is the inventory graph, $\mathbf{V}$ is the node set, $\mathbf{E}$ is the edge set, and $\mathbf{X}$ is the node-feature matrix. In the simulated setting, each shelf-state window covers 30 minutes of warehouse activity. The node features include item count, tag hit rate, visual occupancy, weight change, order event frequency, batch age, storage-zone code, and missingness flag. The graph is updated by sliding windows so that temporary relocation and delayed scanning can be represented as time-local states rather than permanent stock errors.

$$x_i = [q_i, l_i, b_i, s_i, m_i] \tag{2}$$

where x_i is the state of node i, q_i is the quantity feature, l_i is the location feature, b_i is the batch feature, s_i is the sensing feature, and m_i is the missingness flag. The model does not remove missing records before training. Instead, missingness is retained as a signal because a missing scan, delayed inbound confirmation, or absent camera frame may itself indicate an unstable inventory state. Figure 1 shows how shelf nodes, item nodes, sensor nodes, and warehouse events are organized into the inventory graph before cross-domain fusion is applied.

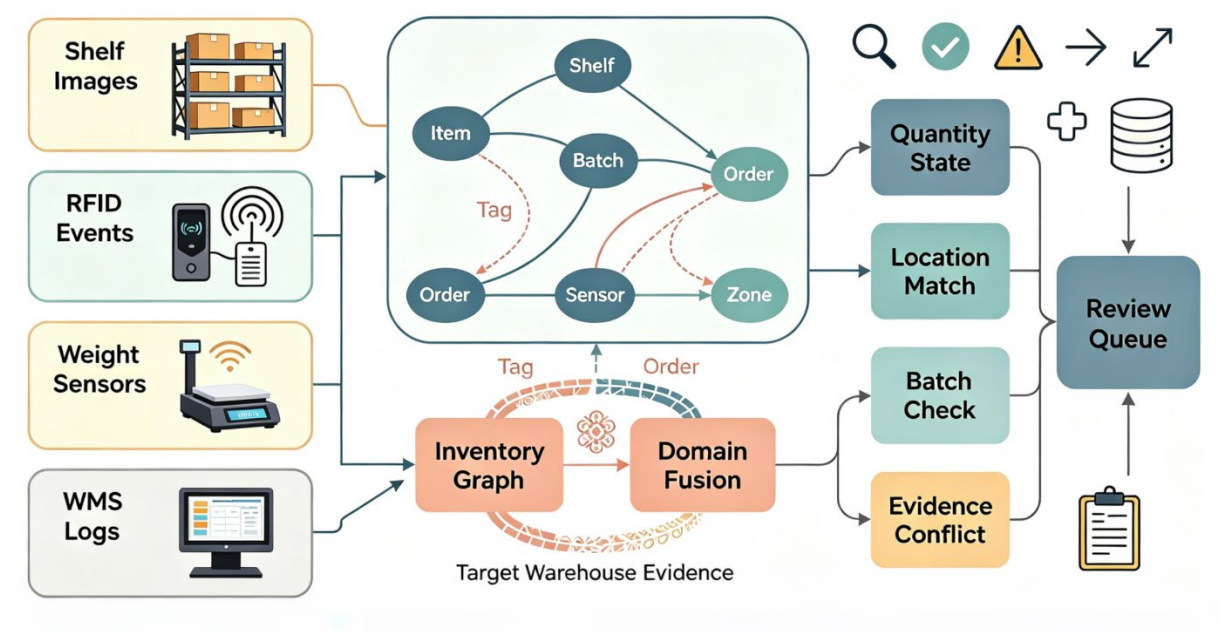


Figure 1. Smart warehouse inventory graph construction and domain fusion workflow.

3.2. Cross-Domain Fusion and GraphGPS Encoding

The cross-domain fusion module separates transferable inventory structure from warehouse-specific background. Source warehouses reliably show ordinary stocking, picking, replenishment, temporary holding, and cycle-count correction, while target warehouses provide fewer labels but contain valuable local event streams. The module aligns common representations and preserves target-specific layout and sensing patterns. This design is necessary because cold-chain warehouses, e-commerce fulfilment centers, spare-parts stores, and mixed manual-automation warehouses can share stock-count logic while differing greatly in turnover speed and sensing reliability. In the proposed model, transferable structure mainly comes from item-shelf compatibility, order-event continuity, and repeated replenishment behavior, whereas target structure reflects operation duration, shelf density, and local sensor coverage.

$$d = \|\mu^s - \mu^t\|^2 \tag{3}$$

where $\mathbf{d}$ is the domain discrepancy, $\boldsymbol{\mu}^s$ is the source-domain feature mean, and $\boldsymbol{\mu}^t$ is the target-domain feature mean. The expression is intentionally short so that the method remains readable in the manuscript. During training, the model reduces the discrepancy of transferable graph features but avoids collapsing all target-domain evidence into the source-domain pattern.

$$h^f = \alpha h^s + (1 - \alpha) h^t \tag{4}$$

where $\mathbf{h^f}$ is the fused representation, $\mathbf{h^s}$ is the source-informed state, $\mathbf{h^t}$ is the target-domain state, and $\boldsymbol{\alpha}$ is the fusion coefficient. The coefficient is learned from warehouse type, sensor coverage, turnover intensity, and recent reconstruction stability. In dense storage areas, α decreases so that target evidence receives more weight; in stable storage areas, source-domain examples provide stronger guidance.

$$H = GPS(G, X) \tag{5}$$

where H is the graph representation generated by GraphGPS. The encoder uses local message passing to model shelf adjacency and shared sensors, and it uses global positional encoding to capture aisle-level and zone-level dependencies. This hybrid structure helps the model identify whether an abnormal count belongs to a single shelf, a nearby cluster, or a wider picking-route issue. Figure 2 presents the GraphGPS-VAE architecture, including domain fusion, graph encoding, latent reconstruction, residual scoring, and review-level output.

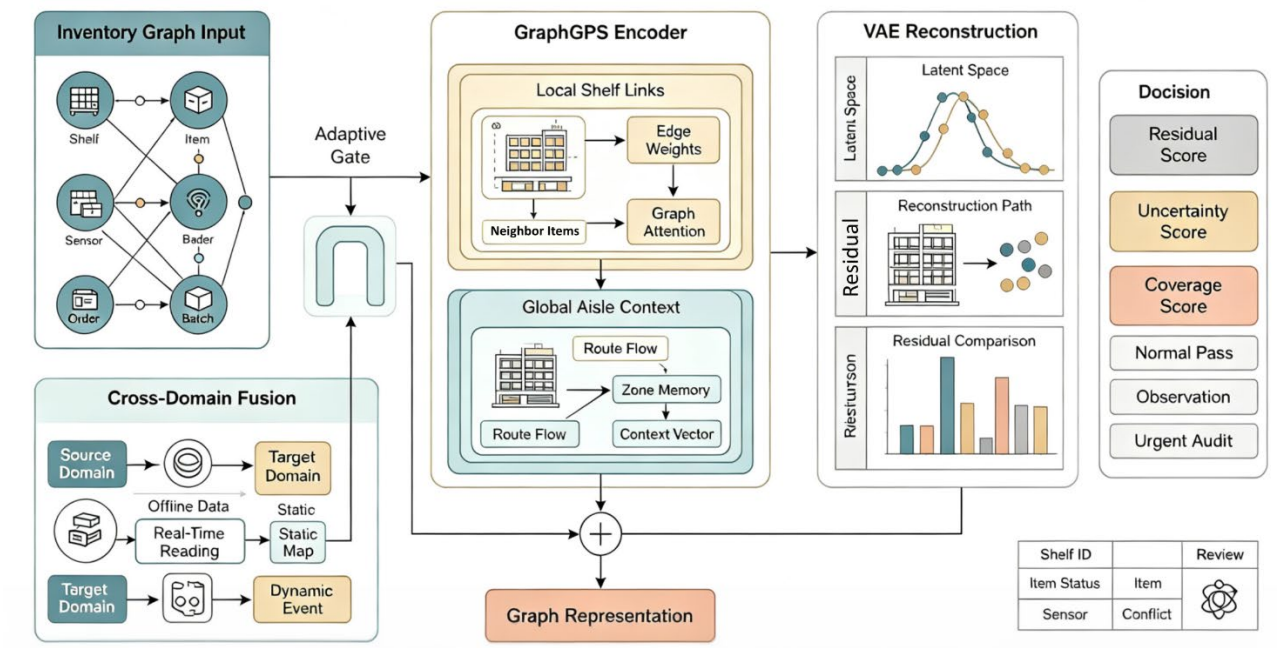


Figure 2. GraphGPS-VAE inventory state reconstruction architecture.

3.3. VAE Reconstruction and Inventory-State Output

The VAE module maps the GraphGPS representation into a latent distribution and reconstructs expected inventory states. The reconstructed state includes expected quantity, expected shelf location, batch consistency, and sensor agreement. The comparison between observed and reconstructed states becomes the basis for review-level warning generation. If a shelf has low visual occupancy, normal weight, and valid tag events, the model may classify the window as uncertain evidence rather than true shortage. If multiple sources agree that an item disappeared without a corresponding picking event, the warning score increases. This reconstruction design is useful because many warehouse problems are contextual contradictions rather than simple classes.

$$z = \mu + \sigma \odot \epsilon \tag{6}$$

where z is the latent inventory state, μ is the latent mean, σ is the latent scale, and ε is a random vector. This step allows the model to represent uncertainty in unstable warehouse windows. For example, a shelf shortly after replenishment may contain temporary evidence mismatch, and the model should not mark every mismatch as a confirmed stock error.

$$r = |q - \hat{q}| \tag{7}$$

where r is the quantity reconstruction residual, q is the observed quantity, and $\hat{q}$ is the reconstructed quantity. The same principle is applied to location consistency and batch consistency. Larger residuals indicate that the current shelf state is difficult to explain under learned normal inventory behavior.

$$S = \beta r + (1 - \beta)u \tag{8}$$

where S is the final review score, r is the reconstruction residual, u is the uncertainty score, and β is set to 0.68 after validation. The output is divided into normal pass, observation, planned review, and urgent inventory audit. The model records the responsible shelf, item, sensor source, and event window so that warehouse staff can check the physical evidence before correcting the stock account.

4. Experiments and Results

4.1. Dataset and Experimental Setting

The suggested multi-source sensing and cross-domain transfer methodology is tested using a simulated smart warehouse dataset. 14 warehouses, 132 storage zones, 48,000 shelf-state windows, 1.92 million sensing events, 21,600 order-operation records, and six inventory-state labels are all included in the collection. Radio-frequency identification tags, visible shelf occupancy, weight-sensor changes, warehouse management logs, transfer records, and cycle-count confirmations are all integrated into each shelf-state window. The experimental configuration for inventory state reconstruction is displayed in Table 1. Both stable-operation and unstable-operation phases are purposefully included in the data design. Unstable times include replenishment peaks, temporary holding, delayed posting, tag overlap, and

shelf movement; stable periods are for regular picking and routine storage. As a result, the assessment will be more in line with actual warehouse management, and the most challenging windows will typically arise during process changes rather than idle storage. Some of the data sets additionally include operator scan omissions, short sensor failures, transient aisle congestion, and shelf-code modifications to prevent an unrealistically perfect simulation. Unless there has been a change in the final inventory status, these occurrences are not regarded as direct abnormal indications; otherwise, they may be innocuous or suspicious in real operation and only truly harmful if they are connected to later stock movements.

Table 1. Smart warehouse inventory dataset configuration.

Item	Value	Domain Role	Main Evidence	Evaluation Use
Warehouses	14	Source/target	Layout, process	Transfer test
Storage zones	132	Zone level	Aisle, shelf type	Graph relation
State windows	48,000	Training sample	30 min window	Reconstruction
Sensing events	1.92 M	Multi-source	RFID, vision, weight	Fusion
Operation logs	21,600	Process record	Inbound/outbound	Traceability
State labels	6	Risk class	Normal and abnormal	Warning output

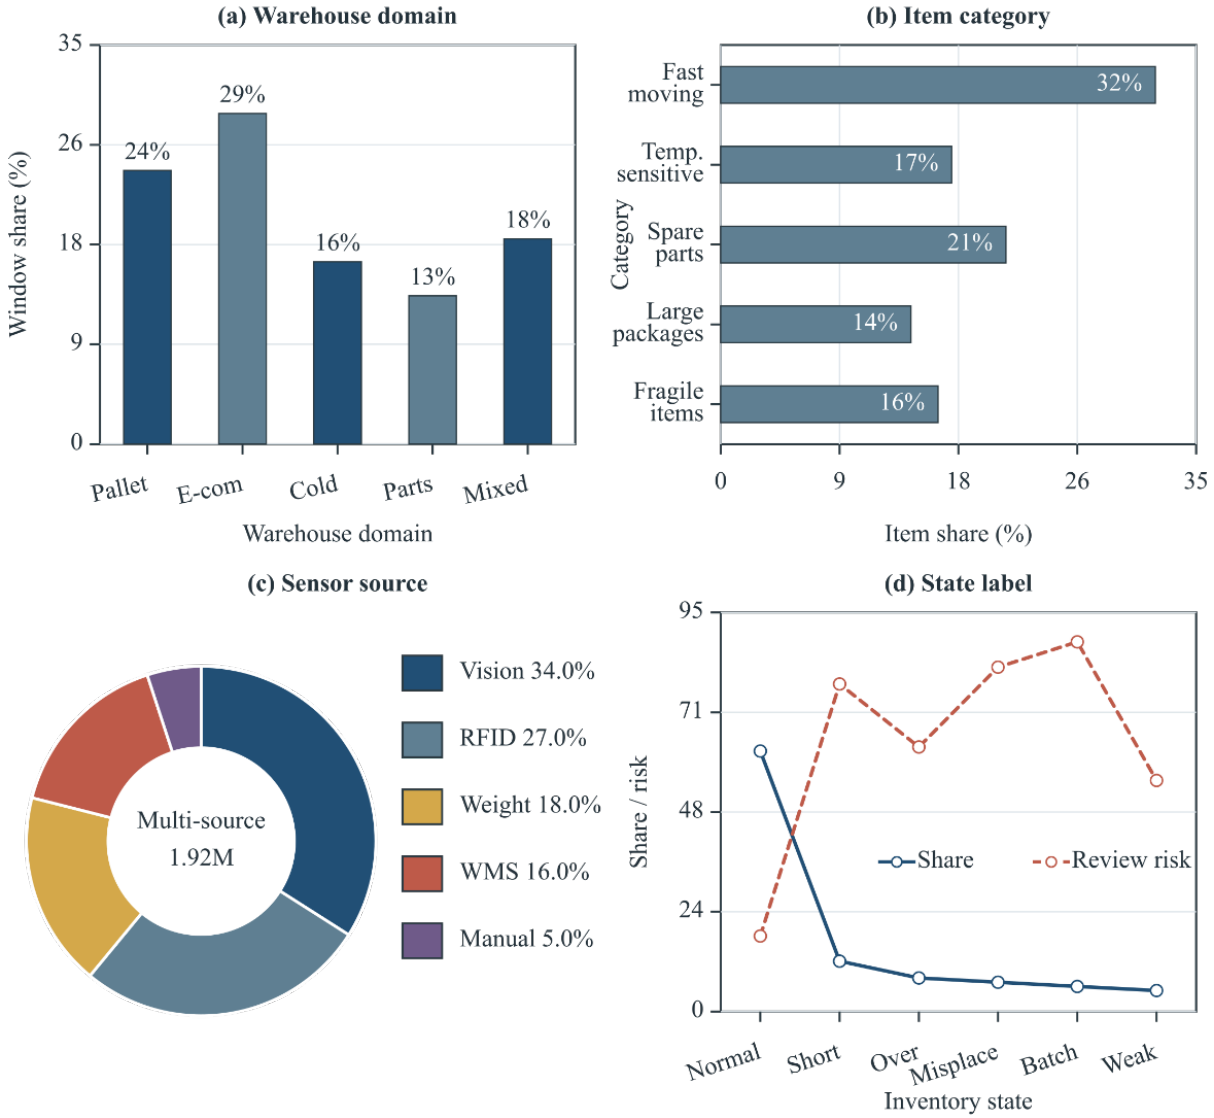


Figure 3. Dataset composition and inventory-state distribution. (a) warehouse domain. (b) item category. (c) sensor source. (d) state label.

Standard pallet storage, e-commerce picking, cold chain storage, spare parts storage, and mixed manual-automation warehouses are the different kinds of warehouses. Normal status, shortage, overstock, misplacement, batch mismatch,

and inconclusive evidence are the labels. The inventory-state distribution and dataset composition are displayed in Figure 3. The warehouse-domain share is shown in Figure 3(a), and because of their frequent shelf changes and extensive operation logs, e-commerce picking and pallet storage have produced the most windows. This section of the picture illustrates that there is more than one warehouse plan in the dataset. Fast-moving products, temperature-sensitive goods, spare parts, large packages, and delicate items are among the item-category proportions displayed in Figure 3(b). Spare parts have longer storage times and increased batch-age uncertainty, while fast-moving items lead to more picking-related mismatches. The relative amounts of visual frames, weight records, warehouse management logs, and radio-frequency identification are displayed in Figure 3(c). The opinion indicates that the fusion might not be trustworthy because no single source covers all the facts. The distribution of inventory-state labels is shown in Figure 3(d); batch mismatch and ambiguous evidence are comparatively uncommon but operationally relevant, whereas normal samples make up the majority. The label imbalance is deliberate; just a few instances require cautious handling, while the majority of warehouses contain numerous accurate records.

Rather than using randomly mixed windows, the evaluation is separated by warehouse domains. Ten warehouses are used for source-domain training, two warehouses are used for validation, and two warehouses are reserved for unseen-domain testing. This setup avoids placing nearly identical adjacent windows from the same shelf into both training and testing sets. Darban et al. [21] have shown that anomaly-detection evaluation should consider temporal dependence and distribution shift, not only randomly mixed samples. In warehouse data, neighboring time periods often preserve the same SKU, quantity, operating mode, and sensor condition. Domain splitting therefore provides a harder and more deployment-oriented assessment because the model must reconstruct inventory states in a facility with a different layout, turnover rhythm, and sensor coverage.

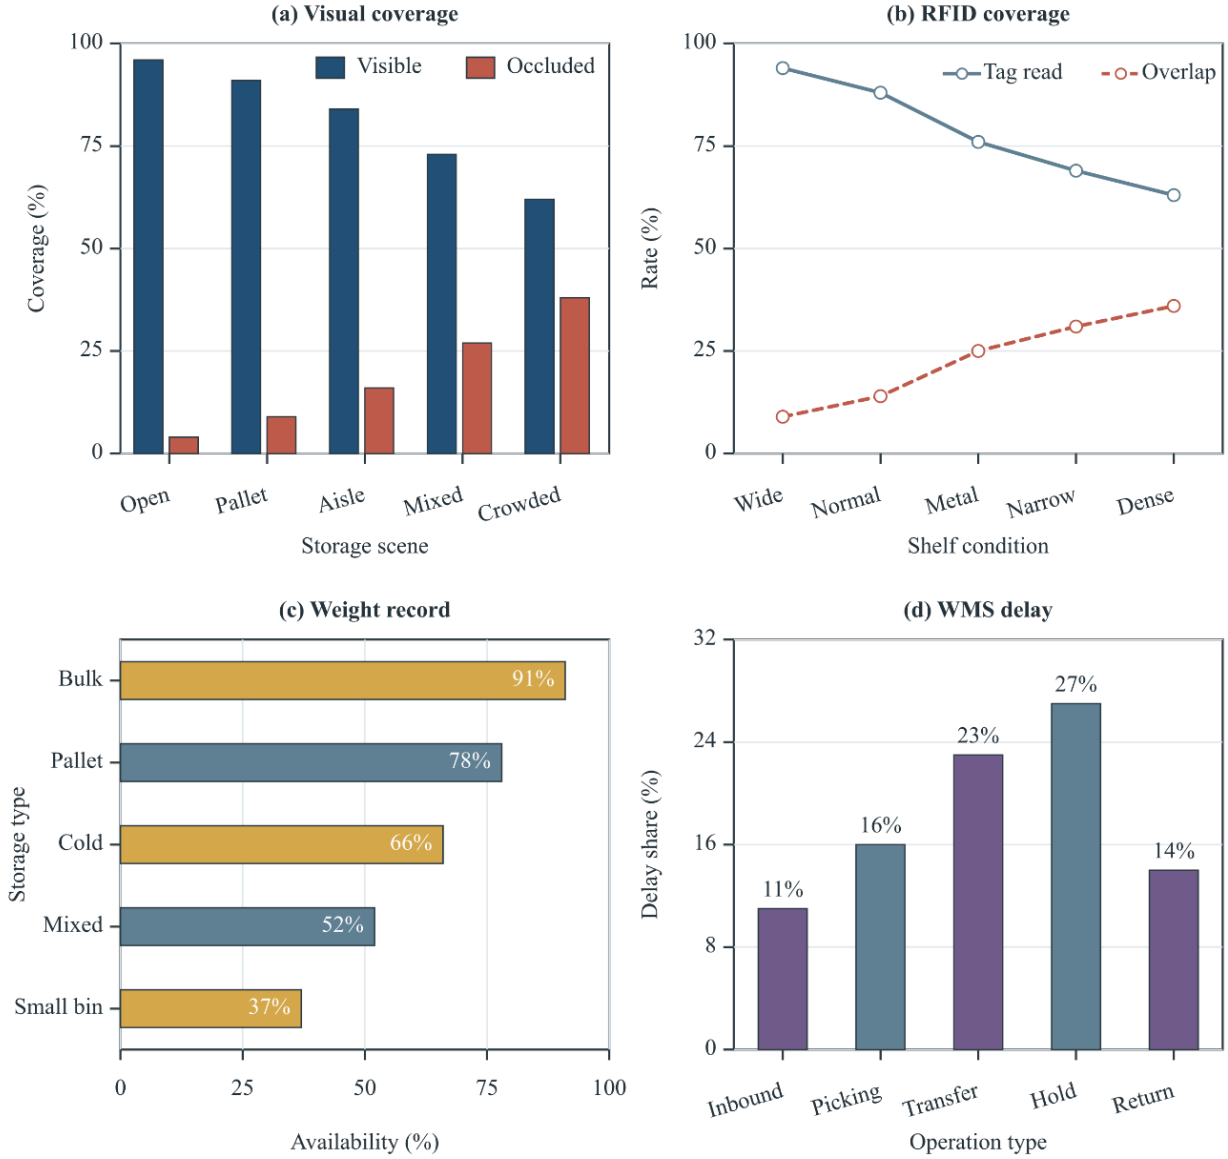


Figure 4. Multi-source sensing coverage and missingness. (a) visual coverage. (b) RFID coverage. (c) weight record. (d) WMS delay.

The sensing completeness and missingness requirements are depicted in Figure 4. Visual frames are comparatively complete in open pallet regions but less so in crowded mixed-shelf areas, as Figure 4(a) illustrates. Uneven aisle illumination, temporary carts, and piled products are the major causes of the drop. The radio-frequency identification coverage is depicted in Figure 4(b), and tag overlap is especially severe in metallic shelf regions and narrow aisles. This finding clarifies why the precise shelf position cannot be verified by tag presence alone. The weight-sensor availability is shown in Figure 4(c); it is low in small-item bins but high in bulk storage. Although weight records are not typically stored in low-value locations, they are highly useful for quantity reconstruction. The delayed warehouse-management activities, specifically batch transfer and temporary holding, are depicted in Figure 4(d). These perspectives demonstrate the instability of single-source inventory models: each sensor has a distinct failure mechanism, and the reconstruction model must decide if the discrepancy is the result of an actual stock exception or merely insufficient evidence. By connecting each disturbance to an actual source of operational uncertainty, the picture also provides a foundation for the subsequent robustness test. Furthermore, the four perspectives demonstrate that the missingness is systematic rather than random. Weight missingness is connected to storage value; log delay is connected to process time; visual missingness is spatially concentrated; and tag uncertainty is connected to shelf material and aisle width. For this reason, rather than eliminating incomplete windows, the suggested model retains missingness as a feature.

4.2. Reconstruction Accuracy and Baseline Comparison

The best overall reconstruction performance is achieved by the suggested GraphGPS-VAE model. The graph convolutional reconstruction baseline is 2.74 units, the quantity means absolute error drops to 2.18 units, VAE without domain fusion is 2.91 units, isolated visual recognition is 3.36 units, and rule-based checking is 4.12 units. Figure 5: Baseline comparison of reconstruction performance. The amount mistake is displayed in Figure 5(a). Because the suggested model combines order events with sensing evidence, it is more stable for commodities that move quickly. After prolonged picking, the isolated visual model is weakened by stacked packages and label occlusion, and the rule-based baseline has a substantial inaccuracy because it regards delayed posting as a direct shortage. GraphGPS-VAE uses shelf adjacency and aisle-level context to achieve 92.7% shelf-location accuracy Figure 5(b). Although the graph convolutional baseline enhances local consistency, it is less appropriate for a transfer job involving long-distance cargo transit. As a result, rather of using individual sensor readings, the integrated warehouse data will be used to determine the inventory condition. The remaining amount mistakes are concentrated in mixed bins and high-turnover shelves, according to error inspection. These are the locations where the products may have been briefly removed, scanned later, or replaced with packages that looked identical. Although many hard mistakes have been transformed into reviewable warning situations, the suggested paradigm does not completely eliminate ambiguity.

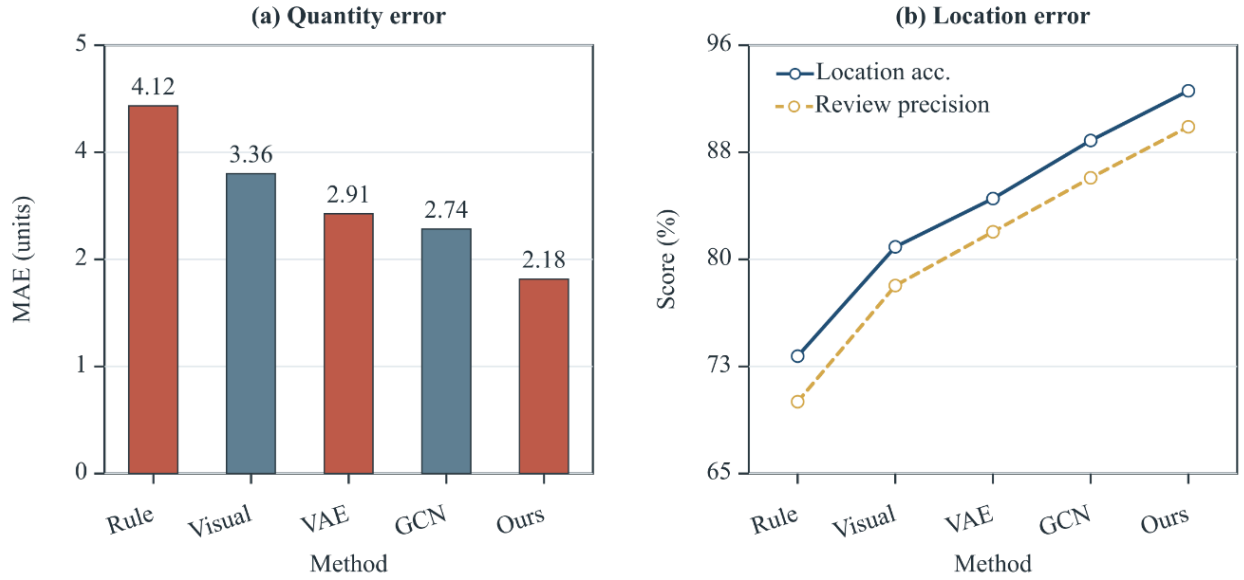


Figure 5. Reconstruction performance comparison. (a) quantity error. (b) location error.

Timely abnormal-state recall is essential because missed inventory exceptions can propagate into replenishment, allocation, and shipment decisions. The proposed model obtains abnormal-state recall of 0.884 and review precision of 0.901. Shortage and misplacement are identified more reliably than batch mismatch because they leave stronger evidence in quantity, tag location, and picking logs. Recent inventory anomaly research [22] supports multi-source warning logic under conflicting evidence. Batch mismatch remains difficult because an item may be physically present while being assigned to the wrong inbound lot. By avoiding frequent escalation of temporary post-picking inconsistencies, the method reduces manual review load by 21.6% compared with rule-based verification.

Cross-domain transfer is shown in Figure 6. Figure 6(a) presents source-to-target reconstruction accuracy after warehouse-layout changes; when transferring from pallet storage to mixed storage, the proposed model loses only 3.1 percentage points, while the graph convolutional baseline loses 6.8 points. Figure 6(b) shows the domain gap under different sensor configurations, with the largest gap appearing when the target warehouse lacks weight data. Figure 6(c) reports adaptation gain after cross-domain fusion, especially for cold-chain and high-turnover warehouses whose operating rhythms differ from standard pallet storage. Figure 6(d) shows unseen-warehouse testing, where the proposed method maintains stronger abnormal-state recall than the baselines. These results indicate that adaptive fusion improves generalization without erasing target-domain evidence.

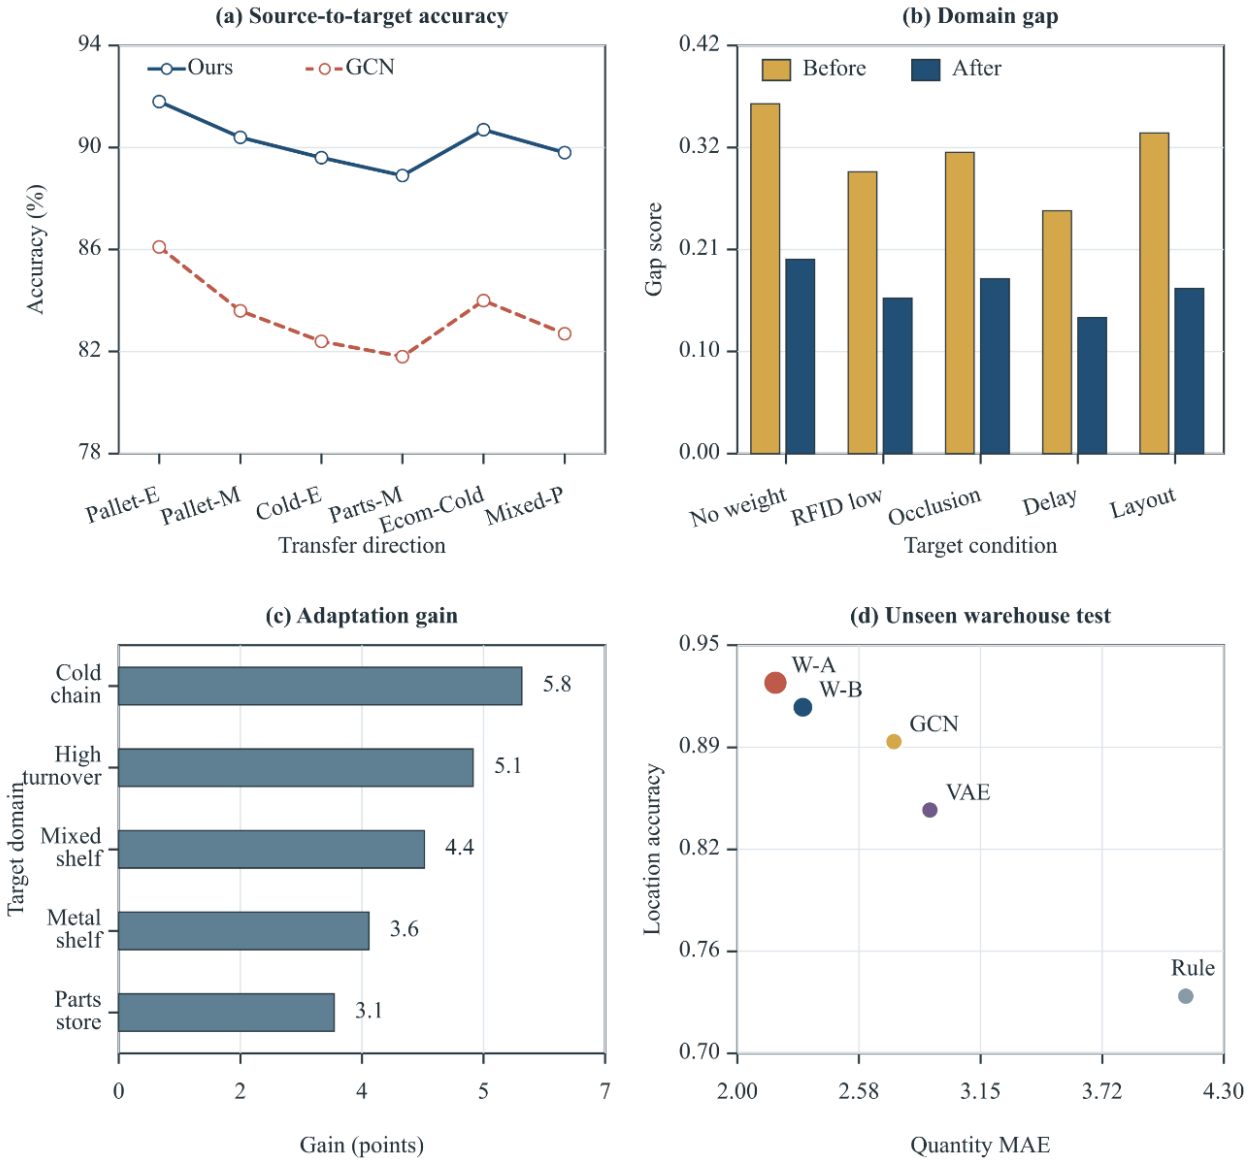


Figure 6. Cross-domain transfer results. (a) source-to-target accuracy. (b) domain gap. (c) adaptation gain. (d) unseen warehouse test.

The ablation analysis clarifies the contribution of each component. Removing GraphGPS reduces location-consistency accuracy from 92.7% to 88.5% because long-range aisle and zone dependencies are not sufficiently captured by standard graph convolution. Removing the VAE reconstruction module lowers abnormal-state recall from 0.884 to 0.842, confirming that reconstruction residuals are useful when abnormal labels are rare. Removing cross-domain fusion increases target-warehouse quantity error from 2.18 to 2.86 units. Studies on variational reconstruction [23] and graph anomaly detection [24] support the value of relational structure and uncertainty-aware residuals. The errors are also different: without GraphGPS, misplacement errors increase; without VAE, ambiguous evidence is forced into normal or abnormal categories; without fusion, the target warehouse inherits source-domain assumptions.

4.3. Operational Warning and Ablation Analysis

Four perspectives of the component ablation findings are shown in Figure 7. Figure 7(a) shows the model without GraphGPS, where adjacent-zone propagation remains weak even though local shelf faults can still be detected. Figure 7(b) shows the model without VAE, where low-frequency batch mismatches are more easily missed despite strong confidence for frequent classes. Figure 7(c) shows static fusion, which performs acceptably for source-like warehouses but poorly when target-domain sensor coverage changes. Figure 7(d) shows the full model, where adaptive fusion, graph encoding, and variational reconstruction jointly reduce error. Recent cross-domain learning research [25] emphasizes the need to retain target-domain residual evidence in different deployment environments. The ablation views therefore represent complementary functions rather than interchangeable modules.

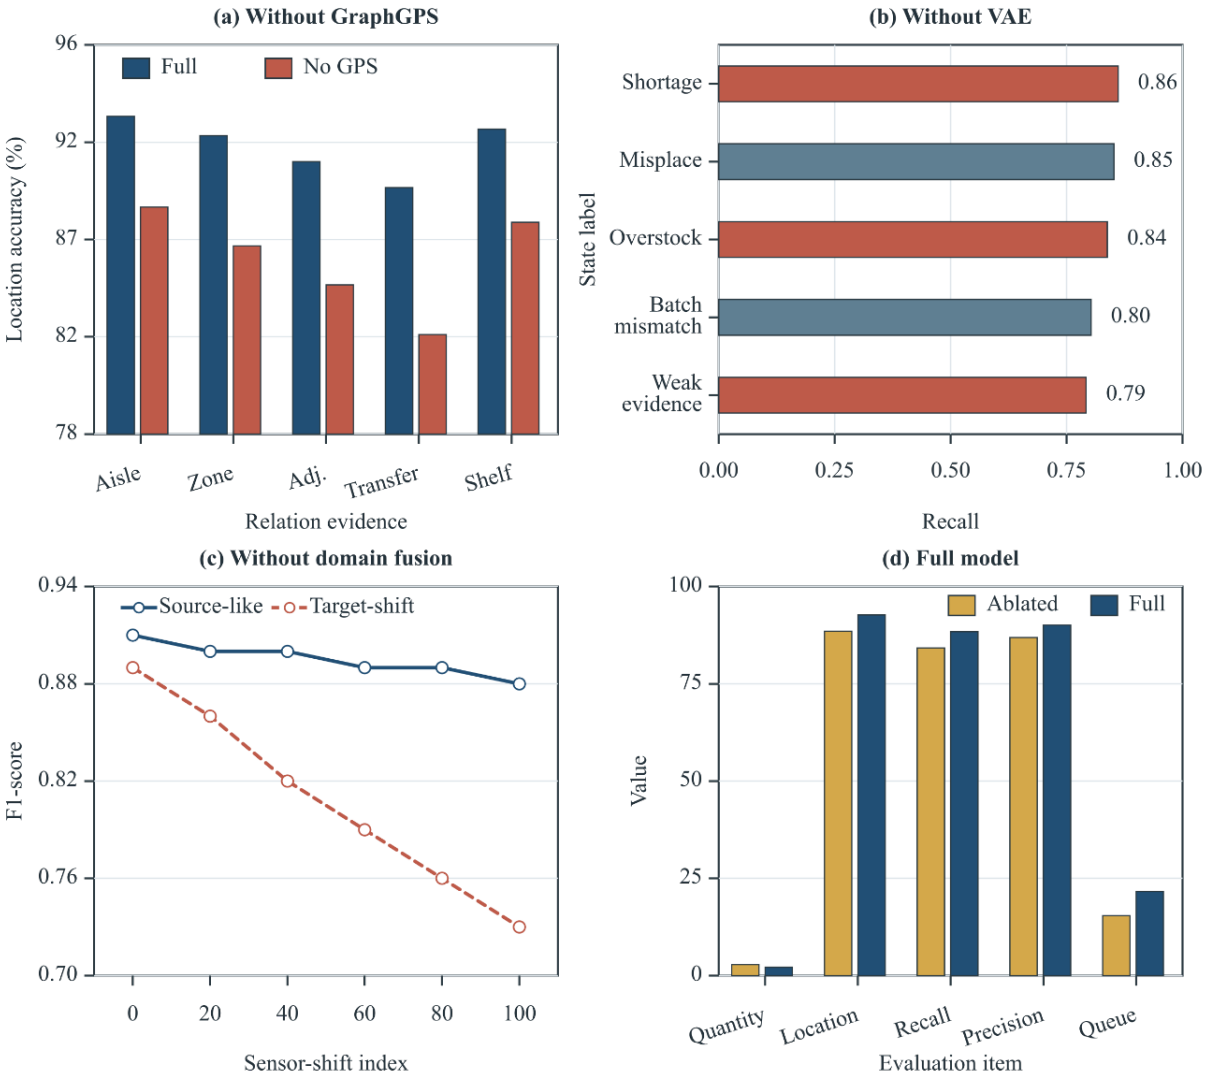


Figure 7. Ablation study of model components. (a) without GraphGPS. (b) without VAE. (c) without domain fusion. (d) full model.

The warning output does not automatically adjust inventory; it is designed for warehouse staff review. A normal pass indicates agreement between digital and physical evidence. Observation means the evidence is weak or temporarily inconsistent, such as a recent picking window without completed posting. Planned review assigns a shelf to the next cycle-count round. Urgent inventory audit indicates a multi-source conflict involving item, location, quantity, or batch status. Figure 8 displays field-oriented reconstruction results. Figure 8(a) maps high-residual bins near suspected shortage, overstock, and batch-mismatch regions, helping staff identify the relevant aisle quickly. Figure 8(b) presents a review-priority queue sorted by warning score, business consequence, and confidence level. The output is therefore organized by action rather than by algorithm confidence alone.

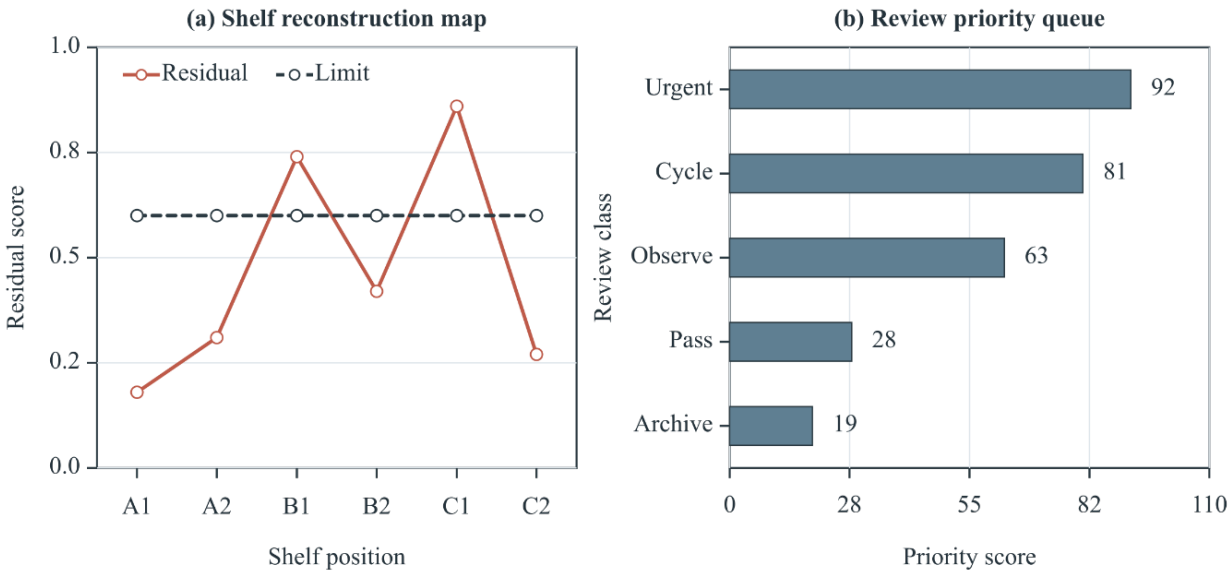


Figure 8. Field-oriented reconstruction output. (a) shelf reconstruction map. (b) review priority queue.

The following are the three ways the model is used in the warehouse. First, the true-shortage-delayed-posting split in the model will make the replenishment assessment more focused. Second, because high-residual shelves may be inspected ahead of typical low-risk shelves, cycle-count scheduling is more effective. Third, because the system has documented the item, shelf, sensor source, and operation window that caused the alert, exception tracing is easier. According to studies on explainable decision-support [26], operational users need evidence channels rather than merely prediction labels. As a result, the model provides a reconstruction card that includes batch evidence, quantity difference, location dispute, missing source, and suggested review action. This card can be linked to warehouse dashboards or portable terminals in an actual deployment. A reviewer can ascertain if the alert is caused by a sensor-quality issue, a delayed posting event, a physical scarcity, or an inaccurate location. As a result, operational rectification is supported by the model's output without hiding the cause of the warning.

Robustness is assessed with four controlled disturbances: visual occlusion, tag-reader overlap, delayed inbound posting, and missing weight records. Under visual occlusion, quantity error increases by 0.34 units, but warning recall is partly preserved through tag and log evidence. In closely spaced shelf areas, tag-reader overlap causes a 2.6-point location-accuracy decrease. Delayed inbound posting shifts more windows from urgent audit to observation because the model recognizes temporary process uncertainty. Missing weight data require cross-domain fusion to retain target uncertainty while borrowing source-domain structure. Sensor-fusion research [27] supports this behavior because a system should remain usable when one source is temporarily unavailable. The limitations remain important: smart logistics studies [28] emphasize process governance, and privacy-aware data-management research [29] reminds us that operational data access must be controlled. The model assumes updated shelf-zone maps and stable sensor identifiers. It may fail under relabelled bins, repackaged items, emergency relocation, widespread WMS outage, fire-safety evacuation, or extended manual bypass. Warnings affecting stock correction, customer allocation, or financial inventory valuation still require human approval [30].

5. Conclusion

This study presented a GraphGPS-VAE model with cross-domain fusion for smart warehouse inventory-state reconstruction. The method constructs a shelf-item-sensor graph from RFID events, visible shelf evidence, weight changes, WMS logs, transfer records, and cycle-count confirmations. GraphGPS models local shelf adjacency and global aisle-level dependencies, while the VAE module reconstructs expected quantity, location consistency, batch relation, and sensor agreement. Compared with rule-based, visual, graph convolutional, and non-fusion VAE baselines on a simulated multi-warehouse dataset, the proposed model reduces quantity error, improves location-consistency accuracy, increases abnormal-state recall, and lowers manual review load. These findings indicate that inventory status is better treated as a relational operational state than as isolated shelf observations.

Several gaps remain. Because the dataset is simulated, uncommon warehouse events such as emergency relocation, widespread sensor failure, mixed-ownership inventory, unusual manual bypass, relabelled bins, and repackaged items are not fully represented. The method also depends on reliable operation logs, updated shelf maps, and stable sensor identifiers. If process records are out of order, cross-domain fusion may propagate false patterns. The current warning score should therefore support inventory monitoring and prioritization, not automatic account adjustment. Cases that influence physical stock correction, customer allocation, financial valuation, or loss investigation require manual verification and process approval.

Future work will validate the model using real warehouse data from different businesses, facility types, and operating routines. Privacy-preserving collaboration should be explored when sites cannot share raw inventory evidence. Robot inspection records, digital twin simulation, staff feedback, and WMS work-order linkage can further strengthen inventory reconstruction for replenishment planning, cycle-count optimization, loss prevention, and financial stock auditing. The reconstruction output may automatically assign high-risk shelves to the next inspection route, but human reviewers should retain authority over official corrections. This direction would make the system an auditable decision-support layer for smart-warehouse operations.

6. Declarations

6.1. Author Contributions

Kristo Pärn contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Gert Berg contributes to validation, analysis, investigation, data collection, draft preparation. All authors have read and agreed with the manuscript before its submission and publication.

6.2. Funding and Acknowledgments

This research received no specific financial support from any funding agency.

6.3. Institutional Review Board Statement

Not applicable.

6.4. Informed Consent Statement

Not applicable.

6.5. Declaration of Competing Interest

Not applicable.

6.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

6.7. Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

7. References

[1] Maheshwari, P., Kamble, S., Kumar, S., Belhadi, A., & Gupta, S. (2024). Digital twin-based warehouse management system: A theoretical toolbox for future research and applications. The International Journal of Logistics Management, 35(4), 1073-1106. https://doi.org/10.1108/ijlm-01-2023-0030

[2] Delgado, F., Garrido, S., & Bezerra, B. S. (2025). Barriers to visibility in supply chains: Challenges and opportunities of artificial intelligence driven by Industry 4.0 technologies. Sustainability, 17(7), 2998. https://doi.org/10.3390/su17072998

[3] Pracucci, A. (2024). Designing digital twin with IoT and AI in warehouse to support optimization and safety in engineer-to-order manufacturing process for prefabricated building products. Applied Sciences, 14(15), 6835. https://doi.org/10.3390/app14156835

[4] Ali, S. S., Marmolejo-Saucedo, J. A., Piedra, R. L., & Weber, G. W. (2026). Integrating digital twins into smart warehousing: A practice-based view framework for identifying and prioritizing critical success factors. Logistics, 10(4), 73. https://doi.org/10.3390/logistics10040073

[5] Villegas-Ch, W., Maldonado Navarro, A., & Sanchez-Viteri, S. (2024). Optimization of inventory management through computer vision and machine learning technologies. Intelligent Systems with Applications, 24, 200438. https://doi.org/10.1016/j.iswa.2024.200438

[6] Culot, G., Podrecca, M., & Nassimbeni, G. (2024). Artificial intelligence in supply chain management: A systematic literature

review of empirical studies and research directions. Computers in Industry, 162, 104132. https://doi.org/10.1016/j.compind.2024.104132

[7] Galkin, A., Samchuk, G., Kopytkov, D., & Thompson, R. G. (2025). Digital twins in logistics: A comprehensive bibliometric analysis for advancing smart cities and sustainable development. Discover Sustainability, 6(1), 853. https://doi.org/10.1007/s43621-025-01754-0

[8] Cai, Y. (2024). Digital twin model construction for intelligent Internet of Things logistics and warehousing systems. Intelligent Decision Technologies, 18(3), 2407-2420. https://doi.org/10.3233/idt-240324

[9] Gao, Y., Zhang, Z., Chen, B., Li, J., Lu, G., Sun, S., & Zong, L. (2025). Time- and frequency-domain fusion for source-free adaptation fault diagnosis. Information Fusion, 117, 102875. https://doi.org/10.1016/j.inffus.2024.102875

[10] Hathaliya, J. J., & Tanwar, S. (2024). Intelligent privacy-preserving data management framework for medicine supply chain system. Security and Privacy, 7(6), e426. https://doi.org/10.1002/spy2.426

[11] Jo, H., & Lee, S. W. (2024). Edge conditional node update graph neural network for multivariate time series anomaly detection. Information Sciences, 679, 121062. https://doi.org/10.1016/j.ins.2024.121062

[12] Zhao, Z., Xiao, Z., & Tao, J. (2024). MSDG: Multi-scale dynamic graph neural network for industrial time series anomaly detection. Sensors, 24(22), 7218. https://doi.org/10.3390/s24227218

[13] Yang, T., Qiao, Y., & Lee, B. (2024). Towards trustworthy cybersecurity operations using Bayesian deep learning to improve uncertainty quantification of anomaly detection. Computers & Security, 144, 103909. https://doi.org/10.1016/j.cose.2024.103909

[14] Porcelli, L., Trovati, M., & Palmieri, F. (2025). Real-time anomaly detection in seasonal time series with conditional variational autoencoder. Applied Soft Computing, 184, 113761. https://doi.org/10.1016/j.asoc.2025.113761

[15] Chu, Z., Xing, S., Han, B., & Wang, J. (2024). A novel cross-domain mechanical fault diagnosis method fusing acoustic and vibration signals by vision transformer. Sensors, 24(16), 5120. https://doi.org/10.3390/s24165120

[16] Song, C., Niu, L., & Lei, M. (2024). A brief survey on graph anomaly detection. Procedia Computer Science, 242, 1263-1270. https://doi.org/10.1016/j.procs.2024.08.145

[17] Alam, M. M., Ugli, R. A. N., Tanha, K. J., & Jun, T. (2026). AI-enhanced digital twin systems for warehouse logistics optimization: A review of challenges with solutions and future directions. ICT Express, 12(2), 459-479. https://doi.org/10.1016/j.icte.2026.01.009

[18] Onsu, M. A., Simsek, M., Fobert, M., & Kantarci, B. (2025). Intelligent multi-sensor fusion and anomaly detection in vehicles via deep learning. Internet of Things, 31, 101561. https://doi.org/10.1016/j.iot.2025.101561

[19] Dereci, U., & Tuzkaya, G. (2024). An explainable artificial intelligence model for predictive maintenance and spare parts optimization. Supply Chain Analytics, 8, 100078. https://doi.org/10.1016/j.sca.2024.100078

[20] Kisten, M., Ezugwu, A. E. S., & Olusanya, M. O. (2024). Explainable artificial intelligence model for predictive maintenance in smart agricultural facilities. IEEE Access, 12, 24348-24367. https://doi.org/10.1109/access.2024.3365586

[21] Darban, Z. Z., Webb, G. I., Pan, S., Aggarwal, C., & Salehi, M. (2025). Deep learning for time series anomaly detection: A survey. ACM Computing Surveys, 57(1), 1-42. https://doi.org/10.1145/3691338

[22] Wang, C., & Liu, G. (2024). From anomaly detection to classification with graph attention and transformer for multivariate time series. Advanced Engineering Informatics, 60, 102357. https://doi.org/10.1016/j.aei.2024.102357

[23] Yan, L., Hu, H., Yang, K., Wang, G., Guan, T., & Zheng, Y. (2025). Variational autoencoder based automatic clustering for multivariate time series anomaly detection. In Wireless Artificial Intelligent Computing Systems and Applications: Lecture Notes in Computer Science (pp. 151-162). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-71470-2_13

[24] Choi, S. H., An, D., Lee, I., & Lee, S. (2024). Anomaly detection based on graph convolutional network-variational autoencoder model using time-series vibration and current data. Mathematics, 12(23), 3750. https://doi.org/10.3390/math12233750

[25] Zheng, B., Ming, L., Zeng, K., Zhou, M., Zhang, X., Ye, T., Yang, B., Zhou, X., & Jensen, C. S. (2024). Adversarial graph neural network for multivariate time series anomaly detection. IEEE Transactions on Knowledge and Data Engineering, 36(12), 7612-7626. https://doi.org/10.1109/tkde.2024.3419891

[26] Zhang, X., Wang, G., Chen, Y., Yang, W., & Wang, G. (2025). Inter-layer explainable variational autoencoder model for multivariate time series anomaly detection. Engineering Applications of Artificial Intelligence, 159, 111585. https://doi.org/10.1016/j.engappai.2025.111585

[27] Yang, T., Jiang, X., Li, W., Liu, P., Wang, J., Hao, W., & Yang, Q. (2025). Cloud-edge collaborative data anomaly detection

in industrial sensor networks. PLOS One, 20(6), e0324543. https://doi.org/10.1371/journal.pone.0324543

[28] Baskak, M., & Coskun, I. Y. (2024). Warehouse digital maturity model: Measuring warehouse digitality through processes. International Journal of Logistics Systems and Management, 1(1), 10064234. https://doi.org/10.1504/ijlsm.2024.10064234

[29] Shafiei, M., & Golpayegani, A. H. (2026). A privacy-preserving distributed protocol for secure data sharing in untrusted supply chain. Scientific Reports, 16(1), 22562. https://doi.org/10.1038/s41598-026-50342-3

[30] Sobot, T., Stankovic, V., & Stankovic, L. (2024). Human in the loop active learning for time-series electrical measurement data. Engineering Applications of Artificial Intelligence, 133, 108589. https://doi.org/10.1016/j.engappai.2024.108589