UNIVERSITY OF WEST ATTICA

eISSN:3136-7658

Applied Manufacturing Intelligence Journal

Vol. 5, No. 2, April, 2026



PPO-Lagrangian Model for Paper Roll Edge Defect Classification Based on High-Speed Line-Scan Texture Images

Florentin Botez[1], Marian Kiriță[1,*]

[1] *Faculty of Automatic Control and Computer Engineering, Gheorghe Asachi Technical University of Iași, Iași, 700050, Romania*

Received 06 March 2026; Revised 01 June 2026; Accepted 05 June 2026; Published 08 June 2026

Abstract

Edge defects of paper rolls reduce converting stability, damage winding quality and may impair subsequent printing or coating operations. High-speed line-scan imaging can capture long, continuous edge textures, but the resulting strips are affected by vibration, illumination fluctuation, paper-fibre variation, edge-position drift and motion blur. This paper proposes a PPO-Lagrangian model for paper roll edge defect classification using high-speed line-scan texture images. Edge-response maps, defect-region indicators, image-quality weights and process-state features are integrated into a texture window so that local texture evidence and acquisition reliability are evaluated together. A classification policy is trained to identify edge cracks, burrs, folds, fibre lift-off, contamination and normal texture changes, while a Lagrangian constraint regulates false alarms, missed serious defects and unstable review decisions. The framework returns defect class, confidence score, constraint state and alarm priority. Simulated experiments cover edge cracks, burrs, folds, fibre lift-off, contamination and normal texture changes. The proposed model achieves 93.7% overall classification accuracy, 91.4% severe-defect recall and a lower constraint-violation rate than ordinary convolutional and Transformer baselines. The constrained optimization also reduces unsafe decision patterns by separating high-risk edge defects from weak texture disturbances. These outputs support line-side inspection, review routing and production-level decision traceability for roll-to-roll paper manufacturing.

Keywords: Paper Roll Edge Defect; Line-Scan Texture Image; PPO-Lagrangian Model; Constraint-Aware Classification; Industrial Quality Inspection

1. Introduction

Paper rolls are widely used in packaging, printing and coating processes because continuous roll-to-roll production offers high throughput and stable material supply [1]. Edge defects are critical because a small crack or fold at the roll side may expand during unwinding and cause web breakage [2]. High-speed line-scan cameras are suitable for paper-edge inspection because they capture long texture strips with consistent spatial resolution [3]. Industrial defect classification must also consider false alarms and missed severe defects because both conditions affect production decisions [4]. These requirements make paper roll edge inspection different from ordinary static image classification. A practical system must distinguish real edge damage from natural fibre texture, vibration marks and illumination variation while preserving enough speed for continuous operation. In actual roll-to-roll production, a defect at the side edge may not remain local. It can propagate along the moving web, disturb the tension distribution and produce secondary defects during rewinding or converting. This makes edge classification more important than ordinary appearance checking because the output may influence whether a roll is accepted, slowed down, trimmed or sent to manual review.

Corresponding author: marian.kir@ac.tuiasi.ro

https://doi.org/10.68406/amij.2026.vol5iss2np6:68-80

➢ 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.

Previously, the method for paper inspection was thresholding, hand-crafted texture descriptors, and manual examination of suspicious edge regions. A general threshold function is overly sensitive to changes in the light environment and thus cannot reliably distinguish between fibre lifting and small edge cracks. Handcrafted descriptors can describe texture roughness and local contrast, but they are often not effective for defects with similar grey-level changes but different process risks. Recent research on industrial vision [5] indicates that line-scan images should be interpreted together with acquisition quality and process status. The paper roll may have a light band that is actually a folded edge, a transient vibration mark or a regular fibre pattern at that location and speed of the strip. These methods are also not suitable for line-scan strips with a regular fibre pattern. The paper texture is not uniform; there are natural grains and small dots, etc. If the modification is not modelled properly, the classifier will mistakenly recognise a normal defect of the paper as a burr or contamination.

Deep learning can learn texture features automatically; however, a standalone classifier may still optimise for the incorrect operating point. A model that optimises overall accuracy may fail to detect rare but serious defects because there are insufficient samples of these defects. Reinforcement Learning with Constraints [6] offers a method to optimise the inspection decision under production risk limits. PPO-based optimisation [7] has stable policy updates and can penalise unsafe alarm behaviour via Lagrangian constraints. Recently, quality-control studies have also found that the deployment model needs to report confidence and review priority in addition to the predicted class [8]. Therefore, this paper introduces a PPO-Lagrangian model to connect texture classification with production constraints. The risk of serious defects is relatively high in a small sample. If only a small number of cracks and folds are shown to the model, it will learn a limited visual pattern and fail to detect serious abnormal conditions. Constraint-aware learning is used to reduce this defect by assigning a higher operating cost to unsafe errors.

A PPO-Lagrangian model for paper roll edge defect classification based on high-speed line-scan texture images is developed in this paper. Section 2 reviews paper roll edge inspection, line-scan texture classification and constraint-aware learning for industrial quality control. Section 3 introduces texture-window construction, the PPO-Lagrangian classification framework and the training objective. Section 4 analyzes dataset composition, classification performance, robustness, constraint behaviour and production alarm output. Section 5 summarizes the main findings, deployment limitations and future validation requirements.

2. Related Work

2.1. Paper Roll Edge Defect Inspection

Quality of paper roll edges is related to cutting, winding, slitting and storage conditions. The common edge defects are cracks, burrs, folded edges, delamination, local contamination and fibre lifting. Inspection research [9] has shown that edge defects are difficult to classify due to the irregular nature of natural paper texture and a gradual defect boundary. A fold may be a line of shade, and a crack can be narrow and broken. Therefore, the visual inspection needs to preserve the local texture details and edge-position information of the region without relying on global image statistics. Edge states can also change during storage and transport. A roll that is suitable immediately after slitting will have a noticeable edge problem during compression or handling. Therefore, the inspection results need to record both the extent and type of defect, and a follow-up in-process audit can determine whether the problem occurred during cutting, winding or storage.

The Production Effects of Different Edge Defects are Not the Same. A shallow fibre mark is acceptable; otherwise, a long-edge crack will result in web breakage during unwinding. Studies on paper manufacturing [10] have classified defects according to their degree of harm to the process. Thus, a classification model can distinguish between serious and non-serious problems to issue alarms accordingly. If the model only returns a category and no risk information, the engineers will not know which samples need to be dealt with urgently. The predicted class may need to be used differently in practice for an alarm. A small burr can be observed; a long edge crack may need to be corrected by taking the roll out of production. Therefore, classification confidence and alarm priority have been added to the output of the model.

Edge defects are also relatively rare in some cases. Most of the line-scan strips are normally textured and only have a small defect. Imbalanced industrial inspection research [11] has shown that rare serious samples need to be kept for evaluation. If the difficult samples are not included, the high-accuracy model will not cover the actual production risk. For the inspection of paper rolls, severe edge cracks and folded edges should be given more weight due to their relatively infrequent occurrence but considerable impact.

2.2. High-Speed Line-Scan Texture Image Classification

Line-scan imaging is different from area-scan imaging in that texture is accumulated line by line during product movement. Continuous-wave acquisition is relatively simple, but it is more likely to have errors due to speed fluctuations, vibration and light changes. Line-scan inspection studies [12] have found that the texture window needs

to have acquisition-quality indicators; otherwise, the same material surface will appear different in different transportation conditions. Indicators of paper-roll edge inspection are strip blur, edge-position drift, exposure fluctuation, and local contrast. In a high-speed line, each image window is only a small section of a long roll edge. The model needs to make a local decision, but the defect may be outside the current window. Overlap of windows is thus employed to ensure continuity and reduce the risk of missing defects near the window edge.

Methods for Texture Classification Include Local Binary Patterns, Gabor Filters, Convolutional Neural Networks and Transformer-based Models. Although deep models can learn multi-scale edge patterns, if the collection conditions are associated with defect labels, they may have inadvertently learned production shortcuts. Recently, texture-recognition research [13] has proposed a powerful model that integrates local anomaly response with contextual position. A defect close to the cut end is different from a similar texture mark in the middle. Therefore, an edge-state representation is needed for the current problem. Random fibre variations and organised defects are present in paper-edge texture. A model that is too local cannot find the long cracks, and a model that is too general fails to identify narrow burrs. Therefore, the feature representation should keep both the local high-frequency information and the broad-area edge-state context.

High-speed inspection is also computationally expensive. The model should have low latency and run continuously; manual cropping of each roll is not needed. Lightweight visual recognition studies [14] have shown that a small feature and decision module can be employed to recognise narrow edge-strip targets. However, light classification cannot reduce severe-defect recall. Therefore, a good model should have good inference speed, classification accuracy and operating conditions. The deployed inspection model is to run for a long time without interruption to production. After changing the roll of the grade, a small camera adjustment and edge position drift are introduced in the output. The above constraints can be used to build a small, constraint-aware model instead of a large network.

2.3. Constraint-Aware Learning for Industrial Quality Control

Constraint-aware learning is also applicable when inspection does not aim only to maximise accuracy. False rejection increases material waste and manual review workload, whereas missed severe defects may lead to downstream production failure. Gu et al. [15] proposed soft constrained policy optimisation to reduce unsafe behaviour under operational constraints. This idea is suitable for paper roll inspection because the model must reduce severe-defect misses while keeping the false-alarm rate within an acceptable range. From this perspective, the inspection system is treated as a controlled decision process, in which classification accuracy, severe-defect omission and unnecessary review cost are evaluated together.

PPO is a policy-gradient method that improves optimisation stability through clipped policy updates, and its policy-learning formulation can also be used for decision calibration. Enkhsaikhan and Jo [16] showed that risk-constrained reinforcement learning with an augmented Lagrangian multiplier can balance reward optimisation and constraint satisfaction. In this paper, the policy corresponds to defect classification actions, the reward is assigned to correct classification, and the constraints are defined by false alarms, severe-defect misses and unstable confidence.

Lagrange's Method Converts the Constrained Optimisation Problem into an Unconstrained Penalty Problem. If the constraint is not met, increase the multiplier and have the model reduce the undesirable behaviour. Research on industrial decision support [17] has shown that, given the changes in the operating environment, adaptive penalties are more suitable than fixed thresholds. In paper roll inspection, this means that the model will change its behaviour according to changes in defect severity or image quality during production runs. Fixed penalties are generally not suitable due to varying process conditions in production. A false alarm in a noisy environment may have a different cause than a false alarm during stable acquisition. Based on the current situation of the constraint, an adjustable penalty can be used instead of a fixed one.

3. Methodology

3.1. Line-Scan Texture Window Construction

A continuous line scan is acquired at the beginning of each paper-roll edge section. Overlapping texture windows are used so that a defect near a window boundary is not overlooked. Figure 1 shows the high-speed line-scan paper edge texture acquisition workflow. the line-scan camera, illumination unit, moving paper edge, texture-window extraction, quality assessment and defect-candidate representation are connected before classification. This structure preserves edge continuity while recording the acquisition conditions that may influence texture appearance.

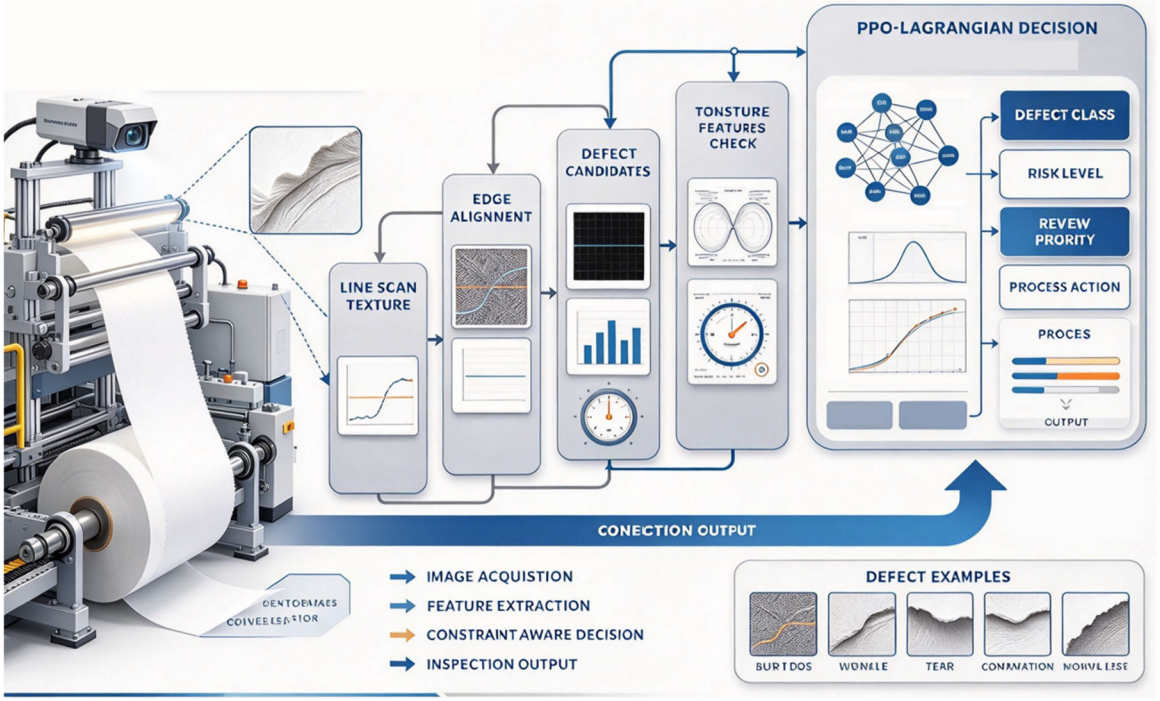


Figure 1. High-Speed Line-Scan Paper Edge Texture Acquisition Workflow

The first feature describes texture abnormality along the edge. It combines local contrast, fibre-direction change and edge-discontinuity response.

$$u_t = [c_t, g_t, e_t, m_t] \tag{1}$$

The observation vector contains contrast, gradient energy, edge offset and motion-quality attributes. It provides the input state for classification, but it is not treated as a direct defect label. The state also includes acquisition information because line-scan images are produced during material motion. A defect-like response during vibration should not be interpreted in the same way as a response obtained under stable acquisition. Therefore, the quality attribute is retained as a model input rather than removed during preprocessing.

The edge response is normalised by local background texture so that natural fibre variation does not dominate the defect signal.

$$r_t = \frac{|u_t - \mu_t|}{\sigma_t + \epsilon} \tag{2}$$

This response shows a sudden change and reduces the effect of ordinary paper texture. Cracks and burrs that have formed at the local irregularities of the adjacent edge pattern can also be corrected. Normalisation is required to address the different base-texture characteristics of paper grades. Coarse-paper grade naturally has more texture variation than smooth-coated grade. Using local statistics, the model compares the candidate region with its own neighbourhood and reduces the impact of grade-specific backgrounds.

3.2. PPO-Lagrangian Defect Classification Framework

Figure 2 shows the PPO-Lagrangian classification framework for paper roll edge defects. The left side receives line-scan texture windows and process-state indicators. The central policy network predicts defect categories, and the constraint module evaluates false-alarm risk, missed severe defects and confidence instability. The Lagrangian multiplier then modifies the penalty assigned to unsafe behaviour. classification and constraint evaluation are therefore learned as a coordinated inspection decision rather than as two independent postprocessing steps.

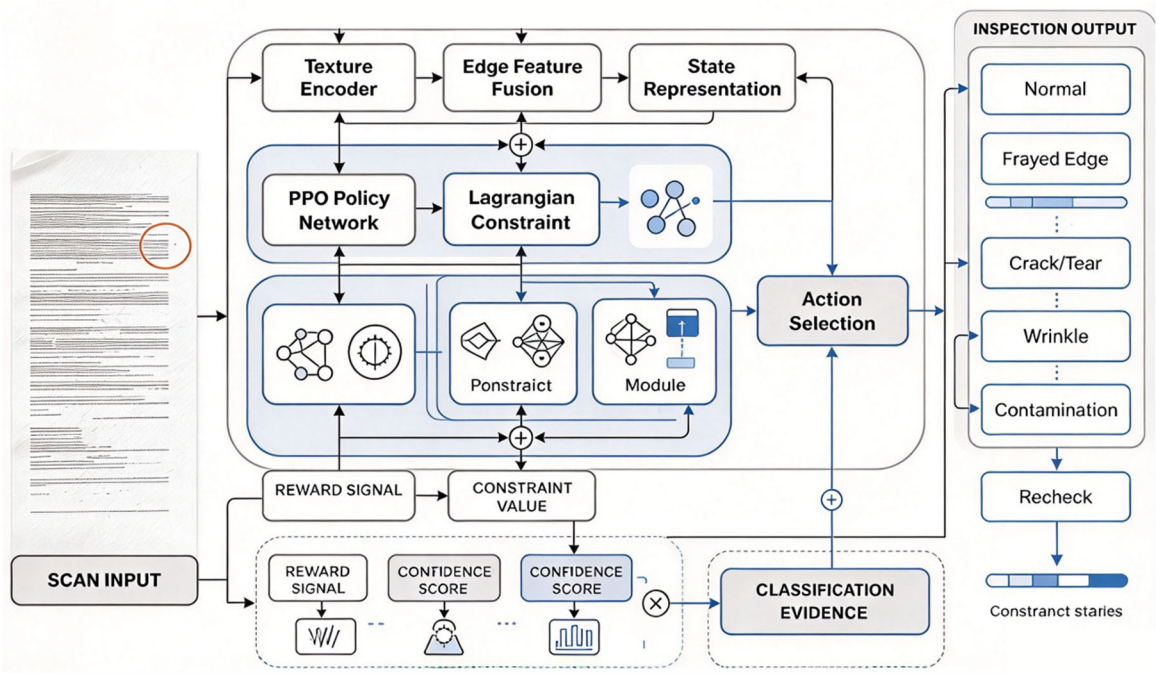


Figure 2. PPO-Lagrangian Classification Framework for Paper Roll Edge Defects

The policy gives the probability of assigning a texture window to each defect category.

$$\pi_\theta(a_t \mid u_t) = \text{softmax}(Wh_t) \tag{3}$$

The policy output is a defect-class decision trained with both reward and constraint information. The action set includes normal_texture, crack, burr, fold, fibre_lifting and contamination. The model therefore learns category selection together with production operating preferences. When weak but high-risk evidence appears, a conservative policy can increase the probability of a severe-defect action without increasing all alarms indiscriminately.

The PPO objective limits overly large policy updates and stabilises training.

$$L_{ppo} = \min(r_t A_t, \text{clip}(r_t, 1-\delta, 1+\delta)A_t) \tag{4}$$

The clipped objective prevents the model from shifting its decision boundary too strongly because of a small number of severe defects in one batch. This property is important for imbalanced production data. Training stability is required because rare severe defects can otherwise create large gradients and unstable false alarms. PPO limits the update size while allowing the policy to respond to reliable severe-defect evidence.

The Lagrangian term represents the adaptive penalty for violating production constraints.

$$L_{\text{lag}} = L_{\text{ppo}} - \lambda C_t \tag{5}$$

When the constraint value increases, the multiplier penalises the policy more strongly. The model therefore learns to reduce unsafe decisions rather than only improve average accuracy.

3.3. Defect Output and Training Objective

The final output contains the defect category, confidence score and alarm priority. The classification probability is derived from the policy network.

$$p_t = \pi_\theta(a_t \mid u_t) \tag{6}$$

The confidence score is calibrated using class probability and image-quality evidence.

$$s_t = p_t q_t \tag{7}$$

The final objective combines classification loss, PPO policy loss and Lagrangian constraint loss.

$$L = L_{cls} + \alpha L_{ppo} + \beta L_{lag} \tag{8}$$

This goal enables the model to identify defects and limit the risk of bad decisions. A high-confidence severe edge crack will be immediately rejected; a low-confidence weak texture anomaly will be sent for manual examination. The Classification Output is for use in production. Support for automatic pass decisions, alarm sorting and review

assignment. Confidence value is the strength of support for the predicted class based on the evidence, and constraint state shows whether the model is close to the unsafe decision region. Thus, the inspection results do not belong to any of the class tables. It is the set of decisions that contains the class, confidence, constraint condition and recommended action. This package is more suitable for the production audit because different parties need different parts of the output.

4. Results and Discussion

4.1. Dataset Construction and Edge Defect Morphology

A dataset containing 7,200 line-scan texture windows of normal edges, cracks, burrs, folded edges, fibre lifting and contamination has been generated [18]. Figure 3 presents the distribution of paper-edge-defect samples. Figure 3(a) reports defect categories, and Figure 3(b) compares severe and weak defect samples. Figure 3(c) shows edge-position coverage, and Figure 3(d) shows process-state groups under stable speed, vibration and speed fluctuation. The data are intentionally imbalanced because normal texture is more frequent than severe edge damage in production. Ambiguous edge-transition windows are retained to avoid an overly favourable evaluation.

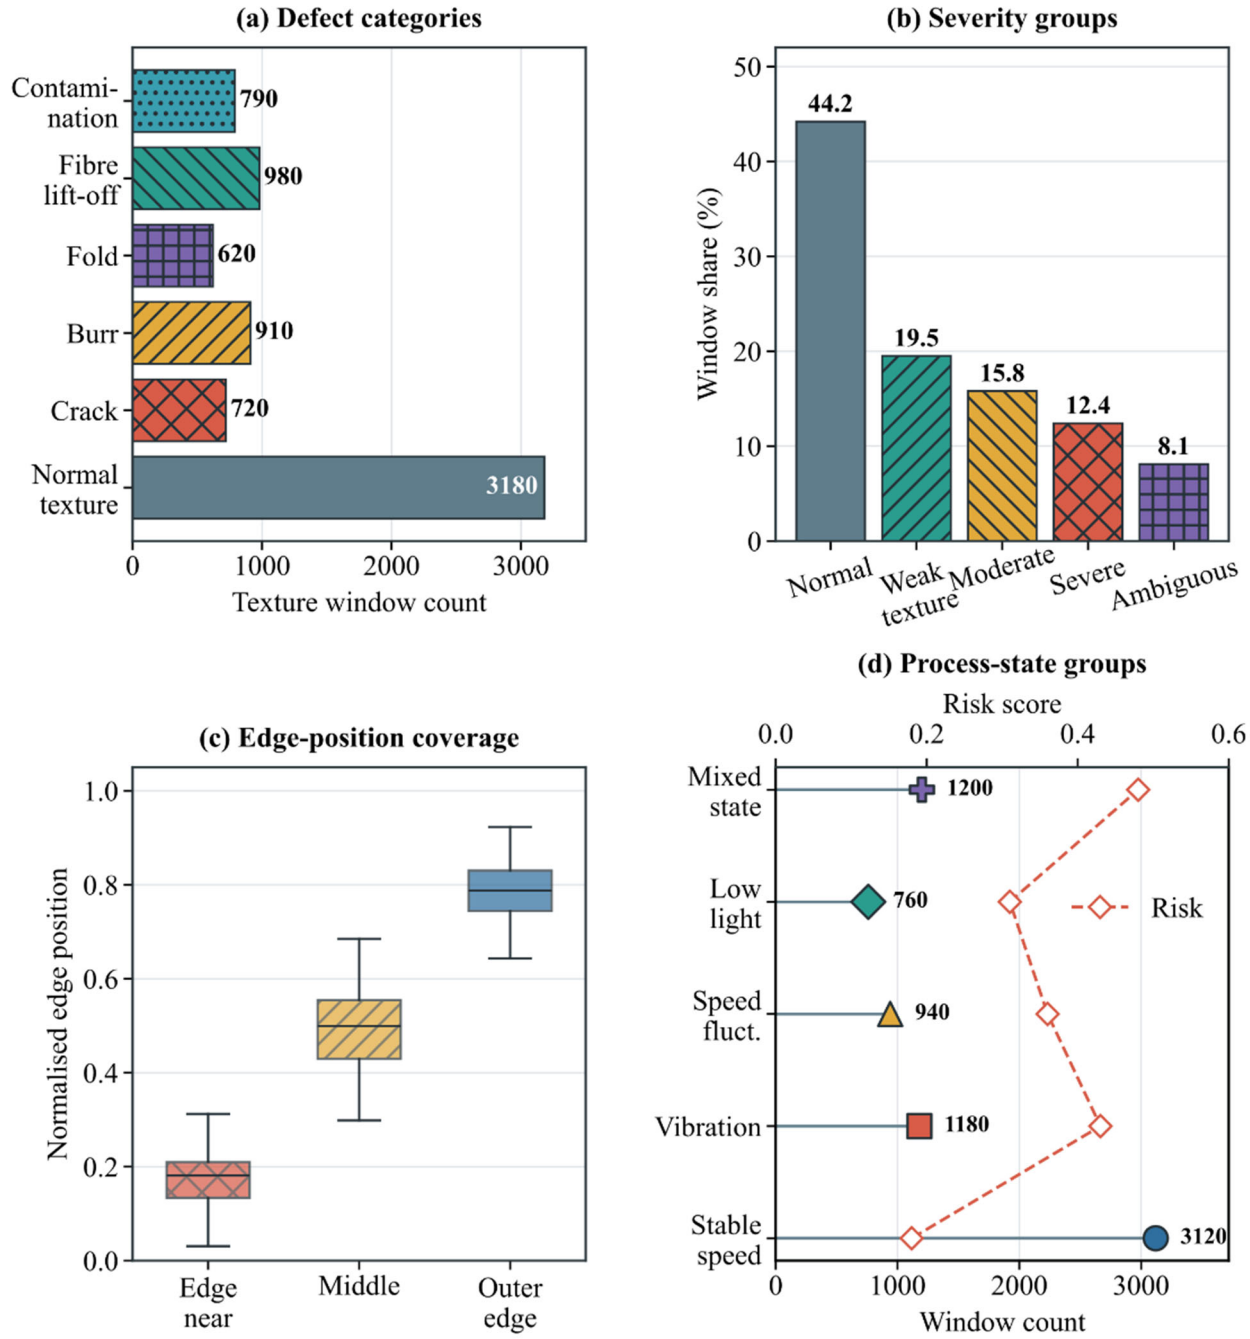


Figure 3. Paper edge defect sample distribution. (a) Defect categories. (b) Severity groups. (c) Edge-position coverage. (d) Process-state groups.

The shapes of defects vary substantially. Edge cracks are often narrow and discontinuous; burrs show rough side protrusions; folded edges form band-like shadows; and fibre lifting produces fine hair-like texture. Research on texture-based defect recognition [19] has shown that multi-scale representation is required for morphological differences. Therefore, the window design preserves local texture detail and edge-position information. A window-level classification objective is more suitable than a single-pixel threshold because burrs, fibre lifting and cracks differ in spatial extent, continuity and production risk.

Figure 4 shows texture variation and line-scan quality statistics. Figure 4(a) presents texture contrast, and Figure 4(b) shows edge-position drift. These two distributions indicate that acquisition conditions can alter the observed edge texture and should not be confused with defect evidence. Quality-aware inspection studies [20] support evaluating image acquisition conditions together with class labels. The feature window therefore retains multiple texture statistics instead of a single intensity value, because folded edges, contamination, vibration marks and normal fibre patterns may change brightness in different ways.

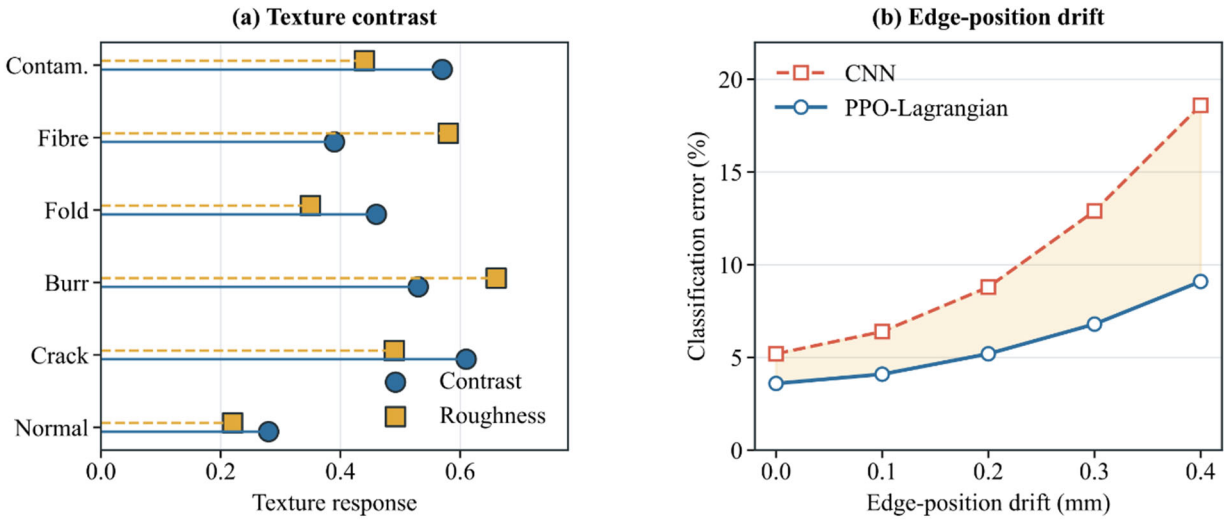


Figure 4. Texture variation and line-scan quality statistics. (a) Texture contrast. (b) Edge-position drift.

4.2. Classification Accuracy and Constraint Performance

A typical supervised classifier [21] is employed to compare the model with CNN, LSTM, Transformer and PPO without restrictions. Figure 5 presents baseline comparison and defect recall. As shown in Figure 5(a), the proposed model achieves 93.7% overall accuracy. Figure 5(b) shows 91.4% severe-defect recall, Figure 5(c) shows lower false alarms than unconstrained PPO, and Figure 5(d) shows more stable confidence under vibration. These subfigures show that constraint-aware optimisation improves production suitability by balancing severe-defect sensitivity, false-alarm control and confidence stability.

Average accuracy does not consider all the inspection results. The CNN baseline is good for general normal textures but cannot recognize some edge cracks. Transformer improves the recognition of extended-area texture but has more false alarms in edge drift. PPO without a Lagrangian constraint improves severe-defect recall but increases false rejections. Constrained learning research [22] suggests that reward-only learning may select aggressive alarms if unsafe actions are not punished. The Lagrangian term is used to prevent constraint violation in motion. Although the model can correctly identify many normal windows, it is still unable to recognise the important severe-crack samples. Even if such a model has good generalisation properties, it is not feasible to apply in practice. As shown in Table 1, the constraint indicators have problems and do not meet the requirements of the plant.

Table 1 shows the accuracy of defect classification and constraint violation. The accuracy, false-alarm rate and constraint violation of the proposed model are all relatively low. The severe-defect missed rate is now 2.9%, and the constraint violation rate is lower than that of PPO without Lagrangian control. Industrial quality evaluation studies [23] have shown that when the cost of various errors is different, these indicators are more representative than accuracy alone. The Difference between PPO and PPO-Lagrangian is particularly important. PPO only learns that severe-defect recall results in a reward and does not sufficiently penalise false alarms. Add a correction force when the alarm behaviour violates the production constraint in the Lagrangian method. Therefore, all of them will be relatively stable.

Table 1. Defect classification accuracy and constraint violation results.

Model	Accuracy (%)	False Alarm (%)	Missed Severe (%)	Constraint Violation
CNN	88.6	5.8	6.7	0.146

LSTM	89.4	5.5	5.9	0.132
Transformer	92.1	6.2	4.1	0.128
PPO	92.8	7.4	3.4	0.153
PPO-Lagrangian	93.7	4.2	2.9	0.081

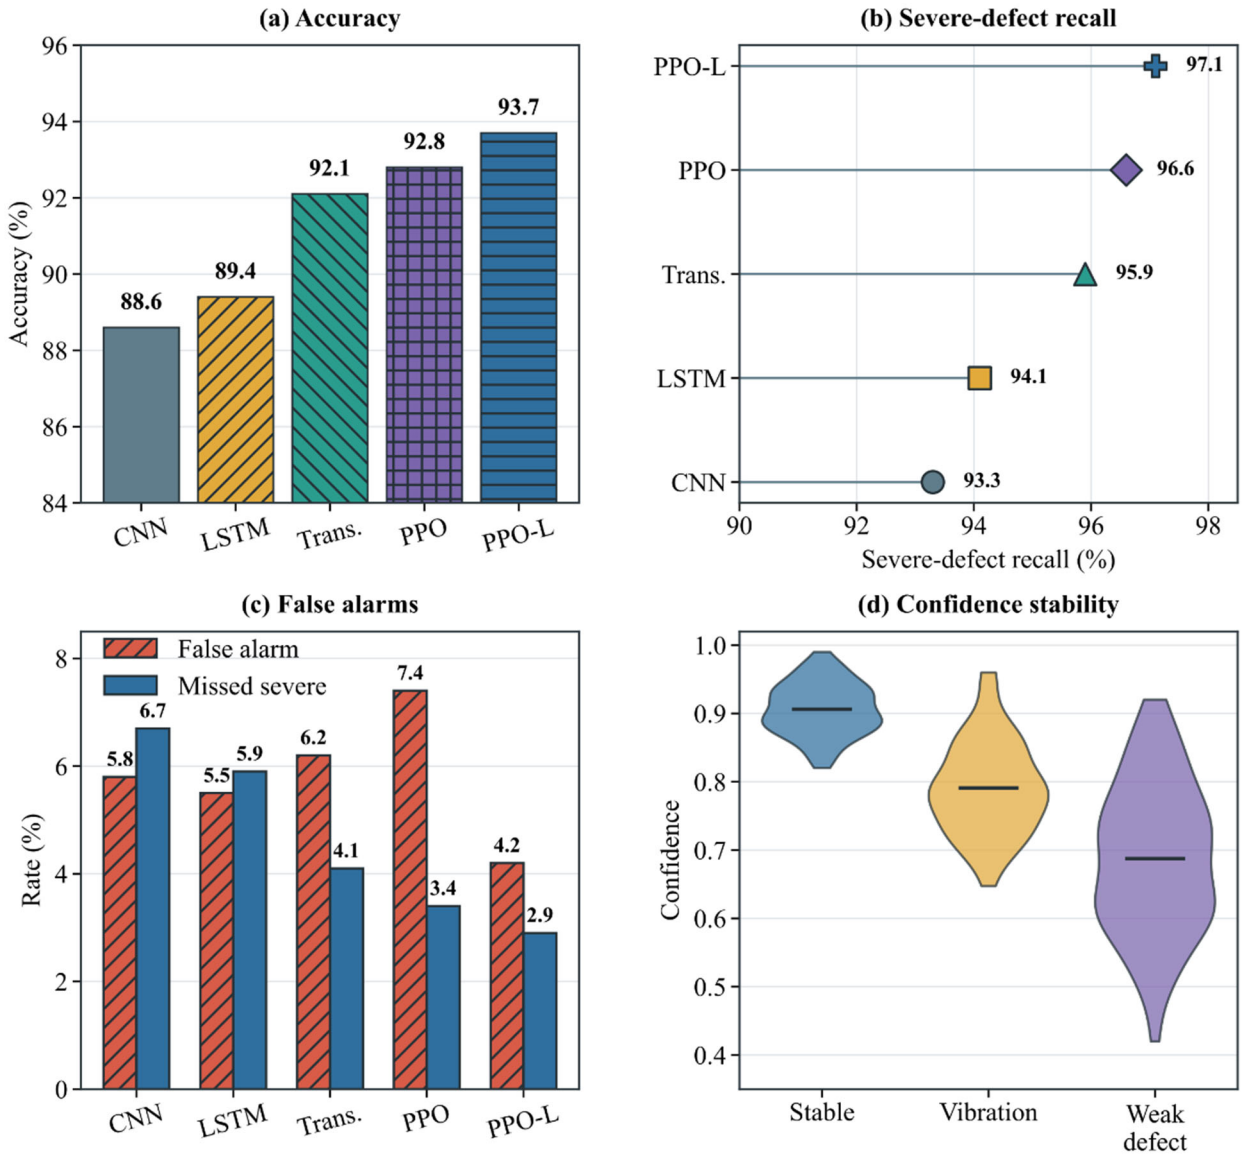


Figure 5. Baseline comparison and defect recall. (a) Accuracy. (b) Severe-defect recall. (c) False alarms. (d) Confidence stability.

4.3. Engineering Review and Production Alarm Output

Figure 6 shows constraint response and confidence analysis. Figure 6(a) presents the change in constraint penalty during training, and Figure 6(b) shows the confidence distribution for severe and weak defects. Research on safe decision-making [24] supports adaptive penalty mechanisms because production risk changes with sample distribution and operating conditions. The constraint curve and confidence distribution should be interpreted together rather than as independent indicators. When penalty is high and confidence is low, unstable illumination, web flicker or mixed labels may be present in the line.

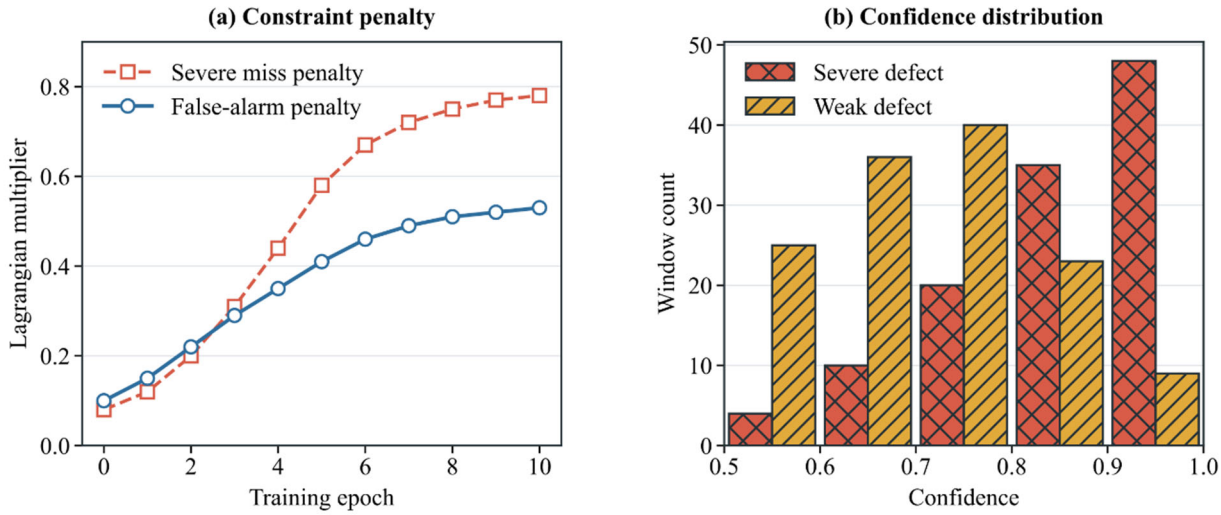


Figure 6. Constraint response and confidence analysis. (a) Constraint penalty. (b) Confidence distribution.

Based on the confidence analysis, weak fibre lifting and contamination are more likely to require manual inspection because their texture patterns are close to normal edge variation. Not all uncertain samples should be rejected by the model. The inspection system should separate urgent alarms from review-priority cases so that the paper production line is not stopped unnecessarily while serious edge damage remains protected against. Low confidence indicates insufficient evidence for an immediate decision, and the sample should therefore be routed to review or sampling inspection.

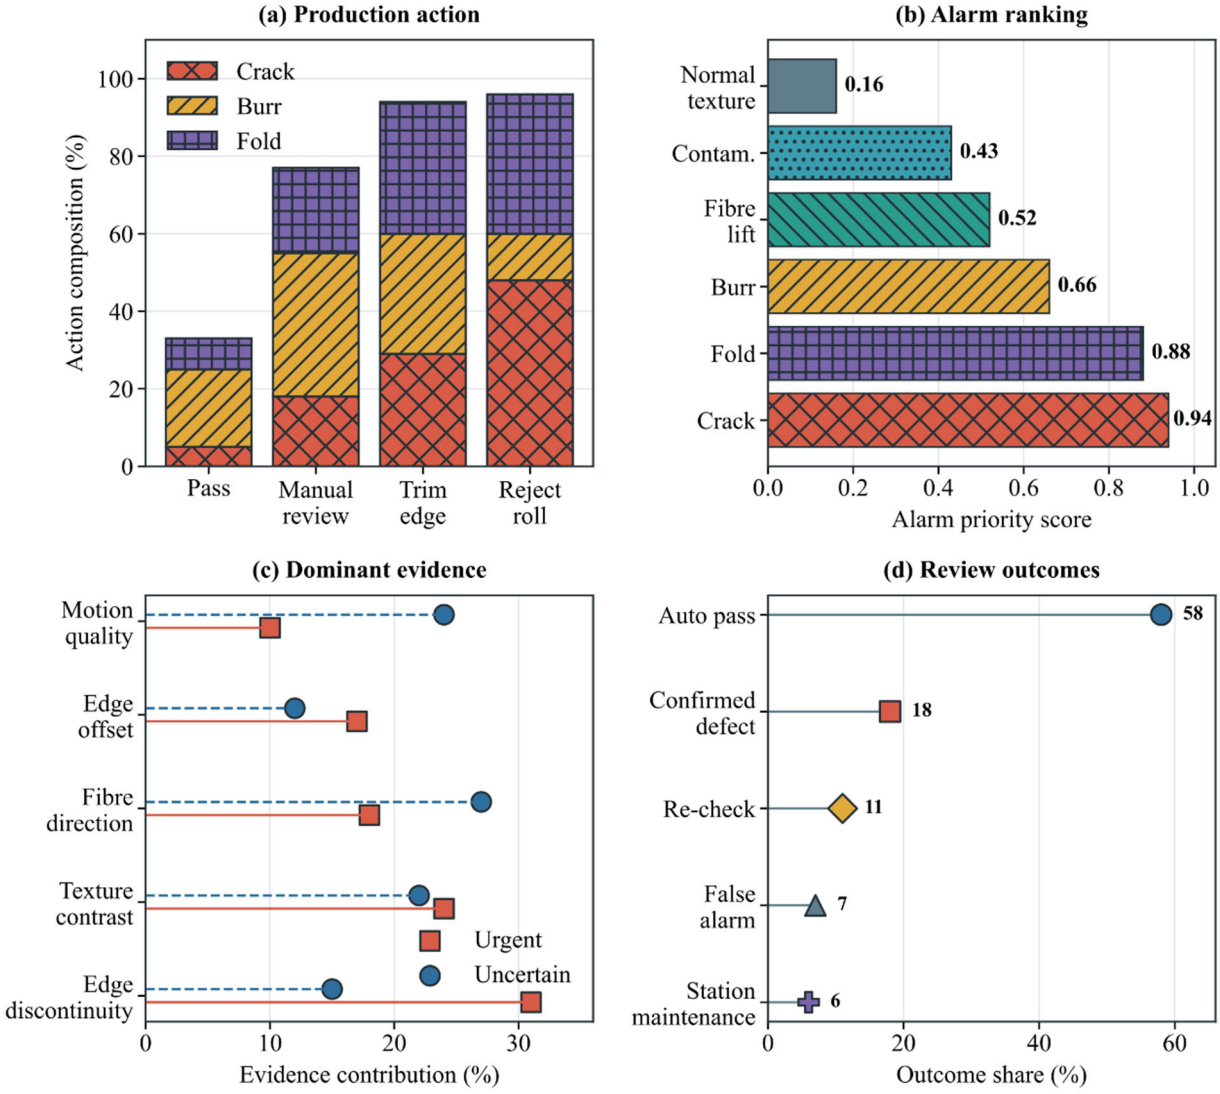


Figure 7. Production review and alarm-priority output. (a) Production action. (b) Alarm ranking. (c) Dominant evidence. (d) Review outcomes.

Figure 7 shows production review and alarm-priority output. Figure 7(a) maps defect categories to production actions; Figure 7(b) records alarm-priority ranking; Figure 7(c) stores dominant texture evidence; and Figure 7(d) summarizes review outcomes. Production decision-support research [25] has found that such records can help engineers distinguish product defects from acquisition disturbances. The subfigures show that rejection, manual review, station maintenance and process diagnosis should be separated. Operators can review whether an alarm is associated with an edge crack, burr, fibre lifting or acquisition disturbance, while maintenance staff can compare repeated alarms with blade replacement, tension adjustment and web-guiding records.

The review log can display process-level attributes. Frequent edge cracks may indicate slitting-blade wear, and repeated fold alarms may indicate unstable winding tension. Predictive Quality Analytics [26] links inspection data with maintenance information and process parameters. The PPO-Lagrangian model supports this connection by outputting class, confidence, constraint state and review priority. Records can be grouped by roll, machine, shift or paper grade, which allows engineers to distinguish mechanical vibration, cutting-tool changes, fibre-lifting patterns and grade-specific texture effects.

Results of robustness to blur, illumination fluctuation and edge-position drift were also obtained [27]. The constrained model reduces false alarms in low-evidence cases and is more suitable for line-scan images with small vibrations. Blur and speed fluctuation reduce the local contrast of narrow cracks, whereas illumination changes alter fibre-texture responses and may be mistaken for contamination or fibre lifting. According to the ablation study, removing the quality-weight term increases the error for speed-fluctuation samples [28]. The texture features, quality weights, PPO policy adjustments and Lagrangian constraint therefore play complementary roles in balancing defect recognition and production risk. Cracks and burrs are localized high-frequency phenomena, whereas folded edges require texture and shape information [29].

The policy output is also an organized response to uncertainty [30]. When both defect probability and constraint state are high, alarm priority is increased; when probability is moderate but image quality is poor, the sample is routed to review or sampling inspection. This ranked output is more applicable than an unstructured alarm list because high-confidence normal windows can be bypassed and questionable edge marks can be retained for manual inspection. The structure can adapt to changes in the production environment better than a fixed-threshold approach [31]. Review-priority output therefore reduces manual workload while preserving uncertain cases that need human confirmation [32].

The current evaluation remains limited because the data are simulated and do not cover all paper grades [33]. Coated paper, recycled fibre and high-gloss surfaces may have different texture statistics. Future work should collect long-term field data covering paper grades, cutting-tool wear conditions and vibration modes. Deployment should also include regular calibration of line-scan illumination and camera alignment [34]. If quality and constraint signals move outside the training distribution, recalibration should be triggered before the error rate increases. Finally, classification output should be connected to downstream web-break and converting records [35], so that high-priority edge alarms can be linked with actual production loss and plant-level cost.

5. Conclusion

This study proposed a PPO-Lagrangian model for paper roll edge defect classification using high-speed line-scan texture images, and the overall conclusion is that edge-response features, acquisition-quality indicators, process-state information and constraint-aware decision logic should be evaluated jointly for traceable roll-to-roll inspection. First, the simulated results show that the proposed model achieves 93.7% overall classification accuracy and 91.4% severe-defect recall, indicating that policy learning with adaptive constraints can preserve sensitivity to high-risk cracks and folds. Second, comparative results show that the proposed model reduces the missed severe-defect rate to 2.9% and the constraint-violation rate to 0.081, while maintaining lower false alarms than unconstrained PPO. Third, the main limitation is that the evaluation relies on simulated texture windows and simplified paper grades; therefore, future validation should include factory line-scan data, paper-grade diversity, cutting-tool wear, vibration modes, camera calibration and downstream web-break or converting-loss records.

The proposed output separates defect class, confidence score, constraint state and alarm priority. This structure supports automatic acceptance for high-confidence normal windows, rejection or trimming review for high-confidence severe defects, and manual review for low-confidence weak texture anomalies. It also allows repeated alarm patterns to be connected with blade wear, tension variation, web guiding or acquisition disturbances. Such separation is necessary because different inspection errors have different production costs and corrective actions.

Future work should acquire extended mill data across paper grades, fibre compositions, coating recipes, slitting tools and line speeds. The model should also be linked with operator feedback, rejected roll length, downstream web-break events and converting-loss records so that alarm severity and constraint bounds can be calibrated against plant-level outcomes. A practical deployment should store image-quality flags, constraint states and final operator decisions

together, enabling engineers to determine whether a quality problem originates from the material edge, the imaging station or the decision threshold.

6. Declarations

6.1. Author Contributions

Marian Kiriță contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Florentin Botez 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. Data Availability

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] Jabbar, K. A. (2024). The effect of employing an optimal web velocity profile on transverse vibrations in roll-to-roll manufacturing. Manufacturing Technology, 24(2), 192–196. https://doi.org/10.21062/mft.2024.034

[2] Lee, Y., Yun, J., Lee, S., & Lee, C. (2024). Image data-centric visual feature selection on roll-to-roll slot-die coating systems for edge wave coating defect detection. Polymers, 16(8), 1156. https://doi.org/10.3390/polym16081156

[3] Kim, B., Shin, M., & Hwang, S. (2024). Design and development of a precision defect detection system based on a line scan camera using deep learning. Applied Sciences, 14(24), 12054. https://doi.org/10.3390/app142412054

[4] Taheri-Kalani, J., Latif-Shabgahi, G., & Shoorehdeli, M. A. (2024). Optimizing multivariate alarm systems: A study on joint false alarm rate, and joint missed alarm rate using linear programming technique. Process Safety and Environmental Protection, 191, 1775–1783. https://doi.org/10.1016/j.psep.2024.09.078

[5] Gao, W., Zhang, J., Liu, K., Li, L., Wang, M., & Ma, B. (2024). Spherical image acquisition and defect detection of nuclear fuel based on the line scan. Measurement, 229, 114371. https://doi.org/10.1016/j.measurement.2024.114371

[6] Bandyopadhyay, S., & Bhasin, S. (2025). Lagrangian-based online safe reinforcement learning for state-constrained systems. Automatica, 179, 112458. https://doi.org/10.1016/j.automatica.2025.112458

[7] Zhou, Y., Tan, H., & Li, J. (2026). Adaptive Lagrangian penalty-enhanced proximal policy optimization for flexible job shop rescheduling with worker workload constraints under concurrent dynamic disturbances. Processes, 14(15), 2519. https://doi.org/10.3390/pr14152519

[8] Thu Trang, T. T., Truong, L. H., & Chien, C. F. (2026). Unified defect classification and unknown defect rejection framework for TFT-LCD array via convolution prototype networks for advanced quality control. Computers & Industrial Engineering, 220, 112293. https://doi.org/10.1016/j.cie.2026.112293

[9] Sakthipriya, G., Padmapriya, N., & Venkateswaran, N. (2024). TSDAnet: Texture strip dual attention network for intraclass texture classification. Signal, Image and Video Processing, 18(11), 7597–7610. https://doi.org/10.1007/s11760-024-03413-9

[10] Wang, C., Wei, X., & Jiang, X. (2024). An automated defect detection method for optimizing industrial quality inspection. Engineering Applications of Artificial Intelligence, 127, 107387. https://doi.org/10.1016/j.engappai.2023.107387

[11] Bai, D., Li, G., Jiang, D., Yun, J., Tao, B., Jiang, G., et al. (2024). Surface defect detection methods for industrial products with imbalanced samples: A review of progress in the 2020s. Engineering Applications of Artificial Intelligence, 130, 107697. https://doi.org/10.1016/j.engappai.2023.107697

[12] Wang, Q. (2025). Monocular line scan vision-based surface defect detection approach for highly reflective bearing balls. Optics Letters, 50(4), 1425. https://doi.org/10.1364/ol.555821

[13] Lainsa, M., & Kim, D. (2024). Zero-shot texture analysis and regression-based deformation recognition for rail anomaly detection. IEEE Access, 12, 138329–138340. https://doi.org/10.1109/access.2024.3466222

[14] Lu, J., Yu, M., & Liu, J. (2024). Lightweight strip steel defect detection algorithm based on improved YOLOv7. Scientific Reports, 14(1), 13267. https://doi.org/10.1038/s41598-024-64080-x

[15] Gu, S., Huang, D., Wen, M., Chen, G., & Knoll, A. (2024). Safe multiagent learning with soft constrained policy optimization in real robot control. IEEE Transactions on Industrial Informatics, 20(9), 10706–10716. https://doi.org/10.1109/tii.2024.3391934

[16] Enkhsaikhan, B., & Jo, O. (2025). Risk-constrained reinforcement learning with augmented Lagrangian multiplier for portfolio optimization. IEEE Transactions on Big Data, 11(5), 2489–2502. https://doi.org/10.1109/tbdata.2025.3533905

[17] Walton, R. B., Ciarallo, F. W., & Champagne, L. E. (2024). A unified digital twin approach incorporating virtual, physical, and prescriptive analytical components to support adaptive real-time decision-making. Computers & Industrial Engineering, 193, 110241. https://doi.org/10.1016/j.cie.2024.110241

[18] Ma, Y., Yin, J., Huang, F., & Li, Q. (2024). Surface defect inspection of industrial products with object detection deep networks: A systematic review. Artificial Intelligence Review, 57(12), 333. https://doi.org/10.1007/s10462-024-10956-3

[19] Jiang, Y., Hu, K., Zhang, X., Zheng, Q., Li, J., & Wu, N. (2024). A coupling enhancement algorithm for ZrO2 ceramic bearing ball surface defect detection based on cartoon-texture decomposition model and multi-scale filtering method. Optics Communications, 554, 130214. https://doi.org/10.1016/j.optcom.2023.130214

[20] Kim, H., Kim, D., Ahn, P., Suh, S., Cho, H., & Kim, J. (2024). ContextMix: A context-aware data augmentation method for industrial visual inspection systems. Engineering Applications of Artificial Intelligence, 131, 107842. https://doi.org/10.1016/j.engappai.2023.107842

[21] Assad, S., Isa, N. A. M., & Saleh, S. A. M. (2026). Hybrid CNN-Transformer models for industrial defect detection: A systematic review. Results in Engineering, 29, 109457. https://doi.org/10.1016/j.rineng.2026.109457

[22] Aggarwal, V., Uddin Mondal, W., & Bai, Q. (2024). Constrained reinforcement learning with average reward objective: Model-based and model-free algorithms. Foundations and Trends in Optimization, 6(4), 193–298. https://doi.org/10.1561/2400000038

[23] Lario, J., Mateos, J., Psarommatis, F., & Ortiz, Á. (2026). A cost model for the investment feasibility of quality inspection technologies in the Zero-Defect Manufacturing era. International Journal of Production Research, 64(17), 7359–7374. https://doi.org/10.1080/00207543.2024.2383780

[24] Zhang, C., Zhang, X., Lan, Y., Gao, H., & Xu, X. (2025). Game-theoretic constrained policy optimization for safe reinforcement learning. IEEE Transactions on Neural Networks and Learning Systems, 36(10), 17990–18004. https://doi.org/10.1109/tnnls.2025.3586603

[25] Bitterling, C., Loske, D., & Klumpp, M. (2025). Improving visual inspection accuracy: Explainable AI for human-centric quality control. IFAC-PapersOnLine, 59(10), 1450–1455. https://doi.org/10.1016/j.ifacol.2025.09.244

[26] 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

[27] Lara, M. A., Hameed, I. A., & Koziorek, J. (2025). Anomaly-focused augmentation method for industrial visual inspection. IEEE Access, 13, 89118–89139. https://doi.org/10.1109/access.2025.3570075

[28] Hao, S., Zhao, J., Ma, X., Fan, S., & Li, T. (2026). HFQE-Det: Frequency-enhanced multi-scale defect detection with confidence calibration for transmission line inspection. Measurement Science and Technology, 37(4), 045403. https://doi.org/10.1088/1361-6501/ae31a6

[29] Guo, J., Yin, M., Guo, X., Wu, X., & Bei, S. (2026). A lightweight surface defect detection network with edge-enhanced

texture reconstruction and dynamic perception. Engineering Research Express, 8(18), 185228. https://doi.org/10.1088/2631-8695/aea74d

[30] Vijayakumar, A., Koilraj, J. A. S., & Rajappa, M. (2025). PharmaNet Deep: Real-time pharmaceutical defect detection using defect-guided feature fusion and uncertainty-driven inspection. International Journal of Computational Intelligence Systems, 18(1), 250. https://doi.org/10.1007/s44196-025-00986-2

[31] Lončarević, Z., Reberšek, S., Šela, S., Skvarč, J., Ude, A., & Gams, A. (2024). Adaptive visual quality inspection based on defect prediction from production parameters. IEEE Access, 12, 93899–93910. https://doi.org/10.1109/access.2024.3424664

[32] Koch, J., Jevremovic, D., & Schüppstuhl, T. (2024). Mobile web app for the digitization and annotation of manual visual inspection tasks. Procedia CIRP, 126, 993–998. https://doi.org/10.1016/j.procir.2024.08.373

[33] Aramkul, S., & Sugunnasil, P. (2026). Multimodal controllable defect synthesis via diffusion-based generation for industrial surface inspection: Bridging data scarcity through material-aware synthetic data generation. IEEE Access, 14, 126680–126705. https://doi.org/10.1109/access.2026.3721567

[34] Shi, B., Wu, F., Wang, C., & Bai, Y. (2024). Methodology and accuracy evaluation of global calibration for multi-line-scan camera system in rail transit tunnel. Measurement, 235, 114914. https://doi.org/10.1016/j.measurement.2024.114914

[35] Wang, F., Chi, X., Wei, L., Song, Y., & Yang, Z. (2024). Multi-position industrial defect inspection using self-training siamese networks with mix strategies. Journal of Industrial Information Integration, 40, 100615. https://doi.org/10.1016/j.jii.2024.100615