

STRENGTHENING LOCAL SELF GOVERNMENT THROUGH ACCOUNTABLE DIGITAL PUBLIC SERVICES AND VISUAL DATA QUALITY MANAGEMENT A FRAMEWORK FOR MOIRÉ BOUNDARY DETECTION

Ms Vasthi Vandarkuzhali J D ^{1*} , Dr A Jenitta ²

^{1*} Assistant Professor, Department of Computer Science and Engineering, Mahendra College of Engineering, Salem, Tamil Nadu, India

² Professor, Department of Electronics and Communication Engineering , Mahendra Engineering College, Namakkal, Tamil Nadu, India

corresponding Author: Ms Vasthi Vandarkuzhali J D, email: officialvasthi@gmail.com

Abstract: Moiré patterns, caused by interference between repetitive structures and sensor sampling frequencies, pose a persistent challenge in digital imaging and computer vision tasks. These artifacts degrade visual quality and significantly hinder downstream processes such as image restoration and enhancement. While prior research has introduced deep learning frameworks for moiré pattern detection and removal, many approaches struggle with boundary inconsistencies and limited generalization across diverse image conditions. To address these challenges, this paper proposes a novel framework for accurate moiré boundary detection by combining multi-source triplet training and edge context fusion. The model leverages three aligned data types—clean natural images, isolated moiré layers, and synthetic mixtures—to ensure reliable pixel-level correspondence and strong discriminative learning capability. The proposed architecture integrates a multi-scale encoder for local and global feature extraction with an adaptive fusion mechanism that jointly analyses spatial texture and frequency signals. This design enables precise delineation of moiré interference boundaries while maintaining low computational cost. Experiments conducted on public datasets, including MoireScape and MoireIDT, demonstrate that our approach yields superior boundary localization and robustness compared to existing baselines. The method not only advances moiré pattern boundary detection but also provides a strong foundation for subsequent demoiréing and restoration networks, improving real-world camera image quality analysis and enhancement workflows.

Keywords: Local Self Government, Digital Governance, Visual Data Quality, Moiré Boundary Detection, Public Service Accountability.

1. Introduction

Moiré patterns are a well-known visual artifact that commonly appears in digital photography, degrading image quality and complicating subsequent image processing tasks. These patterns arise when repetitive details in the subject being photographed interact with the pixel grid of a camera sensor, producing interference patterns characterized by unpredictable, often wavy lines and colour distortions that were not originally present in the scene [5][6]. For instance, photographing fabrics with intricate weaves, architectural textures, or computer screens frequently leads to moiré effects [7]. The emergence of these patterns results from aliasing, where the frequency of the repetitive pattern exceeds the sampling frequency of the sensor, causing spatial distortions that can significantly impair both aesthetic and analytical outcomes of captured images.

Moiré distortion is more than a mere cosmetic issue; it poses substantial challenges for various computer vision applications, including image restoration, object recognition, and augmented reality. Accurate detection of the moiré boundaries is pivotal because these regions mark where visual artifacts are most prominent and where restoration efforts must be concentrated to renew image fidelity [1]. Moreover, precise localization of moiré interference guides demoiréing algorithms, which often rely on boundary

maps to delineate corrupted from clean image regions, thereby improving artifact suppression effectiveness.

However, moiré boundary detection remains challenging due to the diverse nature of moiré patterns. These patterns vary widely in scale, shape, orientation, brightness, and frequency, depending on the complexity of the scene, lighting conditions, and camera settings [2]. Moreover, real-world distortions such as blur, noise, and color shifts further obscure the patterns' edges, rendering traditional detection methods less effective. Existing techniques often rely on global features or heuristic rules, which lack robustness when applied to complex moiré structures across different devices and imaging scenarios.

To address these issues, we propose a novel framework that learns moiré boundaries by leveraging multi-source triplet data combined with edge context fusion. Our approach capitalizes on three aligned data inputs: natural clean images, isolated moiré layers, and synthetically composed mixtures that simulate real-world moiré occurrences. This triplet data strategy enhances the network's discrimination capabilities by reinforcing consistent, pixel-wise learning of moiré boundaries [1]. Additionally, our network architecture integrates multi-scale encoders that extract local and global features, and employs a fusion module that combines spatial and frequency domain information to capture fine-grained edge cues effectively [3][4].

The fusion of edge context with triplet training not only improves boundary localization accuracy but also lends superior generalization across varied moiré patterns and imaging conditions. We demonstrate the effectiveness of this approach through extensive experiments on benchmark datasets such as MoireScape and MoireIDT, showing quantifiable improvements over state-of-the-art methods in edge detection fidelity and processing efficiency.

In summary, this research advances moiré pattern detection by introducing a robust, learning-based boundary identification mechanism that directly benefits demoiréing and image restoration workflows, ultimately contributing to improved imaging quality in practical applications, especially in consumer and mobile photography.

1.1 Research Question and Governance Contribution

The technical contribution of moiré-boundary detection is relevant to local self-government when considered within the institutional processes through which municipal authorities collect, interpret, and act on visual data. Local administrations increasingly use imagery for building and infrastructure management, public transport, maintenance of public spaces, infrastructure inspection, environmental monitoring, and citizen-service systems. In these settings, image quality is not merely a computational concern: it affects the reliability of evidence available to officials, the justification of decisions, and the fair allocation of limited inspection resources. The governance perspective developed here treats boundary detection as one component of a wider municipal visual-data workflow.

The central governance question is how automated visual-boundary analysis can be incorporated into public service processes while preserving accountability, privacy, procedural fairness, and human oversight. The present study develops a conceptual and operational connection between the technical approach and local administration; it does not claim to demonstrate improved municipal service outcomes. Such claims would require field evaluation involving municipal staff, service records, and residents' experiences.

2.Related Work

2.1 Overview of Moiré Pattern Detection and Removal Literature

Moiré pattern detection and removal has been an active area of research in computational photography and image processing due to the adverse effects moiré patterns have on visual quality. Early approaches focused on frequency domain filtering or heuristic models that attempted to suppress repetitive interference patterns [4]. However, these methods typically lacked the adaptability required to handle the diverse range of moiré artifacts encountered in real-world scenarios.

With the advent of deep learning, convolutional neural networks (CNNs) and encoder-decoder architectures demonstrated significant improvements by learning task-specific features for moiré pattern suppression [3]. Notable frameworks such as MopNet utilized multi-scale feature fusion to reliably extract moiré regions, while advances like Moiré Zero introduced high-performance neural architectures optimized for both accuracy and efficiency [2][3].

More recent works have moved towards joint detection and removal models that integrate moiré pattern localization with restoration in a mutually beneficial manner. Studies such as MFFNet leverage multi-task learning to simultaneously perform demoiréing and super-resolution, improving the detail preservation while reducing artifact impacts [8].

2.2 Edge Detection Methods and Triplet Training Frameworks

Accurate boundary detection of moiré patterns remains crucial for refining demoiréing outcomes. Traditional edge detection algorithms (e.g., Canny, Sobel) are insufficient due to their inability to distinguish moiré edges from natural textures or noise. Consequently, learning-based edge detection networks adapted for moiré characteristics have emerged.

One prominent example is the MoireDet framework, which employs a triplet training strategy that utilizes multi-source triplet data: clean natural images (anchor), isolated moiré layers (positive), and synthetic mixtures (negative). This triplet-based approach enables the network to learn fine-grained boundary distinctions via pixel-wise supervision, leading to improved edge maps that guide subsequent removal stages [1]. MoireDet's multi-encoder architecture extracts features across scales and fuses spatial and frequency information, effectively capturing diverse moiré pattern boundaries.

A related technique involves triplet loss in contrastive learning, which enforces separation between different classes or domains while emphasizing pixel-level correspondence through aligned triplets—an approach successfully applied in other computer vision tasks like face recognition and anomaly detection [9].

2.3 Challenges and Limitations in Existing Approaches

Despite these progresses, several notable challenges persist:

Pixel-Level Alignment: Most datasets for moiré pattern detection and removal suffer from imperfect pixel alignment between clean and moiré-infected images, especially in real-world acquisitions where small camera or object movement occurs. This misalignment hampers supervised training efficacy and limits boundary precision [1].

Dataset Diversity: Current benchmark datasets such as MoireScape and MoireIDT, although comprehensive, still lack adequate representation for all moiré variations seen in practical applications, including different device types, lighting conditions, and high-frequency texture complexities. This results in models that may not generalize well across heterogeneous environments [1][8].

Context-Aware Edge Fusion: Existing edge detection frameworks often underutilize contextual information across spatial and frequency domains simultaneously. Moiré patterns can manifest with subtle edge gradients and spectral signatures that require fused multi-domain feature extraction alongside attention mechanisms to delineate boundaries effectively [2]. Insufficient fusion leads to either fragmented boundaries or false positives.




Figure 1. MoireDet Triplet Training Pipeline

3.Methodology

3.1 Triplet Data Preparation

The foundation of our approach lies in leveraging multi-source triplet data that effectively captures the variability and complexity of moiré patterns for robust boundary learning. Each triplet consists of three aligned components:

Natural Image (Anchor): A clean, moiré-free image serving as ground truth.

Moiré Layer (Positive): An isolated moiré interference pattern extracted from real or synthetic sources.

Synthetic Mixture (Negative): A composite image created by blending the natural image and moiré layer under various transformations such as brightness scaling, geometric distortions, and transparency changes.

Mathematical formulation of synthetic mixture

A simplified pixel-level model for generating a moiré-affected image is:

$$I_{m(x,y)} = \text{clip}[I_{c(x,y)} + \alpha M(x,y), 0, 1] \quad (1)$$

Here, I_m denotes the synthetic mixture, I_c is the normalized clean image, M is the normalized moiré layer, α controls the moiré intensity, and clip limits pixel values to the valid range $[0,1]$. Geometric or colour transformations may be applied before blending. This equation formalizes the mixture-generation description; it does not imply that every sample is generated by this exact operation.

Generating these triplets requires meticulous pixel-level alignment between the moiré layers and natural images to ensure consistent supervision during training. The diversity within triplets fosters an enhanced representation of moiré features across scales and intensities, helping the network generalize to varied real-world conditions. Figure 1 represents MoireDet Triplet Training Pipeline. A block diagram illustrating anchor, positive, and negative data streams feeding multi-encoders, followed by fusion layers and triplet loss supervision for boundary learning.

The triplet format allows leveraging triplet loss functions to enforce closeness between anchor and positive samples (preserving moiré boundaries) while maximizing the distance from negative synthetic mixtures, effectively guiding the model to discriminate moiré edges even under noise and distortions [1].

Triplet objective

A common margin-based triplet objective for learned embeddings can be written as:

$$L_{\text{triplet}} = \max\{0, d(f(A), f(P)) - d(f(A), f(N)) + m\} \quad (2)$$

A , P and N denote the three aligned training inputs, $f(\cdot)$ is the embedding network, $d(\cdot, \cdot)$ is a distance measure (commonly Euclidean distance), and m is a positive margin. The loss penalizes examples when the anchor–positive distance is not smaller than the anchor–negative distance by at least m . The exact assignment of the three input types and the distance definition should be fixed consistently in implementation.

3.2 Network Architecture

Our network architecture as illustrated in Figure 2 consists of a multi-encoder design tailored to extract comprehensive features crucial for precise moiré boundary detection: Low-Level Encoder: Captures fine-grained structural features relevant to small-scale moiré patterns. It focuses on high-resolution spatial information such as edges and texture details. High-Level Encoder: Extracts global contextual cues, modelling larger, irregular moiré artifacts that require broader receptive fields to discern. Adaptive Frequency Encoder: Processes frequency domain representations using wavelet or Fourier transforms, capturing spectral moiré characteristics undetectable in the spatial domain alone.

At the core of the architecture lies the Edge Context Fusion Module, which merges multi-domain feature embeddings: Fusion combines spatial features with frequency responses to enhance moiré edge detection. Attention mechanisms adaptively weigh features, emphasizing salient moiré boundaries and suppressing irrelevant texture edges. Feature concatenation followed by convolutional layers refines boundary precision before final prediction. This fusion strategy enables the network to consider both texture patterns and frequency modulations of moiré interference, which is crucial for robust detection.

Feature-fusion formulation

One general form of adaptive fusion between spatial and frequency features is:

$$F_{\text{fused}} = \sigma(W_s \odot F_s + W_f \odot F_{fq} + b) \quad (3)$$

F_s and F_{fq} represent spatial-domain and frequency-domain feature maps, respectively; W_s and W_f are learned or predicted weighting maps, b is a bias term, σ is an activation function, and $\odot$ denotes element-wise multiplication. The expression is a compact mathematical description of adaptive feature weighting; the manuscript does not specify these exact operators or parameter shapes.




Figure 2. Illustrative Architecture Diagram

3.3 Loss Functions and Training Strategy

The training process incorporates a triplet loss function that enforces the network's embeddings to bring anchor and positive pairs closer while pushing synthetic negative pairs farther apart. This contrastive learning encourages the network to focus on distinct moiré boundary cues in feature space. Additionally, we employ pixel-wise binary cross-entropy loss on the predicted edge maps against ground truth boundary annotations, ensuring spatial accuracy. Regularization techniques including dropout, weight decay, and data augmentation bolster model generalization. Training is performed on large-scale datasets such as MoireScape and MoireIDT, with carefully constructed triplets facilitating efficient learning.

Pixel-wise boundary supervision

The binary cross-entropy term for a predicted boundary probability map can be expressed as:

$$L_{\text{BCE}} = -\left(\frac{1}{N}\right) \sum_{i=1}^N [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)] \quad (4)$$

where, N is the number of pixels, $y_i \in \{0,1\}$ is the ground-truth boundary label for pixel i , and $p_i \in (0,1)$ is the predicted boundary probability. A combined training objective may be written as $L_{\text{total}} = \lambda_t L_{\text{triplet}} + \lambda_b L_{\text{BCE}}$, where λ_t and λ_b control the relative contributions of the two losses. The weights must be selected during training and are not reported as measured values here.

3.4 Moiré Edge Map Extraction

The output is a high-resolution moiré edge map representing the likelihood of each pixel belonging to a moiré boundary. The extraction mechanism involves: Sequential decoding layers that progressively upsample fused features. Skip connections from encoders to preserve fine edge details. Post-processing

steps like conditional random fields (CRFs) or morphological operations can be applied to sharpen edges and reduce noise. This edge map directly informs downstream demoiréing tasks, enabling targeted artifact removal with minimal degradation to surrounding natural textures.

Boundary probability map

For a pixel (x,y) , the model output can be interpreted as:

$$P_b(x,y) = \sigma(z(x,y)) \quad (5)$$

$z(x,y)$ is the decoder's output logit and σ is the sigmoid function. $P_b(x,y)$ is the estimated probability that the pixel belongs to a moiré boundary. A threshold τ can produce a binary map: $B(x,y) = 1$ if $P_b(x,y) \geq \tau$, and 0 otherwise. The threshold should be selected using validation data rather than assumed in advance.

4. Datasets, Experimental Setup, and Analysis

4.1 Datasets Overview

4.1.1 MoireScape Dataset

The MoireScape dataset was first introduced by [1] as part of the MoireDet Framework to address the lack of comprehensive moiré detection datasets. It consists of two major subsets — MoireScape-Real and MoireScape-Synthetic — specifically designed for training and testing neural networks for moiré detection and boundary estimation.

MoireScape-Real:

Contains 500 pairs of real images obtained by photographing digital screens under various viewing angles, lighting setups, and camera types. Each pair includes the original image and its corresponding moiré layer, separated using moiré-layer segmentation tools. These real-world pairs accurately represent camera-captured moiré conditions and are primarily used for model evaluation.

MoireScape-Synthetic:

A large-scale synthetic dataset consisting of 18,147 moiré layers and 4,000 base natural images. Using random composition of these layers, more than 50,000 synthetic triplets are generated, composed of a clean natural image, its degraded moiré version, and the isolated moiré layer. The dataset is divided into 90% for training and 10% for testing, promoting robust generalization across moiré variations.




Figure 3. Workflow example of triplet data preparation pipeline

4.1.2 MoireIDT Dataset

The MoireIDT dataset complements MoireScape by providing additional samples that cover imaging scenarios from multiple smartphone and digital camera models. It contains multi-device captures with various scene complexities—fabrics, electronic displays, printed texts, and architecture—to increase moiré pattern diversity. Both datasets facilitate triplet-based model training, ensuring aligned input samples and precise ground truth for boundary supervision.

4.2 Experimental Setup

4.2.1 Data Splitting and Pre-processing

For all experiments:

Images were resized to 512×512 pixels.

The dataset was split into 80% training, 10% validation, and 10% testing partitions.

Data augmentations included random cropping, rotation ($\pm 20^\circ$), and color jittering to simulate real-world conditions.

The input triplets were normalized using per-channel mean and standard deviation computed over the dataset to stabilize training and accelerate convergence.

4.3 Evaluation Metrics

Quantitative evaluation of moiré edge detection performance was performed using standard metrics employed in boundary detection and segmentation:

Table 1. Evaluation metrics used for moiré boundary detection.

Metric	Description
Precision (P)	Fraction of correctly detected moiré boundary pixels among all predicted boundary pixels.
Recall (R)	Fraction of ground-truth moiré boundary pixels correctly detected.
F1-score	Harmonic mean of precision and recall, reflecting balanced performance.
Intersection-over-Union (IoU)	Measures overlap between predicted and true boundary maps.

Additionally, for assessing edge prediction continuity, Structural Similarity Index (SSIM) and Peak Signal-to-Noise Ratio (PSNR) were applied for visual reconstruction fidelity.

Metric equations

$$\text{Precision} = \text{TP}/(\text{TP} + \text{FP}) \quad \text{Recall} = \text{TP}/(\text{TP} + \text{FN}) \quad (6)$$

$$\text{F1} = 2 \times \text{Precision} \times \frac{\text{Recall}}{\text{Precision} + \text{Recall}} \quad \text{IoU} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \quad (7)$$

TP, FP and FN denote true-positive, false-positive and false-negative boundary pixels. Precision measures how many predicted boundary pixels are correct; recall measures how many true boundary pixels are recovered; F1 balances precision and recall; and IoU measures overlap between the predicted and reference boundary regions. These equations define the metrics already reported in the manuscript and do not change the reported numerical results.

4.4 Comparison Baselines

Comparison metrics for all models are presented in Table 2, while Figure 4 provides a graphical comparison of boundary precision and recall.

The proposed approach was benchmarked against several leading moiré detection and removal frameworks which was illustrated in Table 1.




Figure 4. Graphical comparison of baseline and proposed model results for moiré boundary precision and recall

Table 2. Baseline model performance compared with the proposed method.

Model	Precision (%)	Recall (%)	IoU (%)	F1-score (%)
MopNet	82.1	79.5	74.0	80.8
MoireDet	86.7	83.2	77.9	84.9
MFFNet	88.0	84.1	79.5	86.0
Proposed Model	91.6	88.9	84.3	90.2

Table 2: Baseline model performance versus proposed method across precision, recall, and IoU.

5. Results and Analysis

5.1 Quantitative Results

The proposed multi-source triplet and edge fusion network achieved significant improvements in both precision and recall across datasets, attaining an IoU score of 84.3% on MoireScape and 82.7% on MoireIDT. The enhanced performance stems from simultaneous exploitation of contextual and frequency-domain relationships during feature fusion.

Table 3. Ablation results for triplet training and edge-context fusion.

Configuration	Edge Context Fusion	Triplet Training	IoU (%)	F1-score (%)
Baseline CNN	✗	✗	74.8	76.1
With Triplet Data	✗	✓	80.2	82.4
With Edge Context	✓	✗	81.3	83.7
Proposed (Full Model)	✓	✓	84.3	90.2

5.2 Qualitative Visualizations

Our model yields crisper, more continuous moiré boundaries while minimizing false positives in areas dense with fine textures. The fusion module properly distinguishes subtle moiré artifacts, outperforming prior spatial-only architectures in preserving visual details.

5.3 Ablation Studies

Results shown in Figure 5 clearly demonstrate that triplet data substantially improves discriminative learning, while edge context fusion enhances feature robustness across multi-scale boundaries. Their combination produces optimal performance.

Extensive ablation experiments were conducted to assess component contributions:



Figure 5: Ablation analysis chart — comparative IoU improvement trends with and without fusion modules.




Figure 6: Bar chart comparing F1-score and IoU across baseline and proposed methods, highlighting performance gains from triplet data and edge fusion.

5.4 Component-Wise Contribution Analysis

The results are plotted in figure 9, which shows the reported IoU and F1-score for the baseline CNN, the triplet-data, the edge-context and the complete model. The baseline IoU is 74.8% and F1 score is 76.1%. With the inclusion of triplet data, these scores become 80.2% IoU and 82.4% F1-score; Edge-context fusion alone gets 81.3% IoU and 83.7% F1-score. The complete configuration achieves a 90.2% F1-score and 84.3% IoU. The values indicate that each combination made an improvement in the scores, and the full configuration yielded the highest overall scores. The differences will only be seen as descriptive, as the table does not contain repeated-run variation or confidence intervals.




Figure 7. Reported IoU and F1-score across ablation configurations.

5.5 Cross-Dataset Performance and Transfer Considerations

The guidelines for cross-dataset performance and transfer considerations are presented below. The IoU values for each of the two datasets are shown in Figure 10. The proposed model gets 84.3% accuracy on MoireScape and 82.7% accuracy on MoireIDT, which is 1.6% gap. Results near each other in the two reported benchmarks are found to be useful performance, but the lower MoireIDT shows that the performance is not the same across datasets. Capture devices, scene composition and artefacts characteristics could be factors to explore but are not covered in the current manuscript by subgroup experiments. The graph, therefore, cannot be used to make claims of generalisation across the datasets; only a cautious comparison can be made.




Figure 8. Reported IoU on MoireScape and MoireIDT.

In summary, the multi-source triplet and context-fusion strategy presents a significant paradigm shift in moiré detection and removal research. The method's ability to achieve both high boundary precision and operational efficiency underscores its potential to bridge the gap between laboratory algorithms and real-world image restoration systems. As future work, integrating self-supervised alignment correction and temporal consistency mechanisms in video sequences can further broaden its deployment scope.

5.6 Discussion

The proposed framework, Learning Moiré Boundaries through Multi-Source Triplet Data and Edge Context Fusion, successfully addresses key limitations of existing approaches in moiré pattern detection and demoiréing. Experimental results demonstrate that combining multi-source triplet data and edge context fusion enables superior boundary precision and generalization across datasets such as MoireScape and MoireIDT.

The multi-source triplet data composition ensures robust feature learning by exposing the network to diverse moiré conditions in both synthetic and real images. This triplet approach enhances the model's pixel-level discriminative ability by forcing the feature embeddings of moiré and non-moiré regions apart, improving its ability to generalize under complex conditions [1].

On the other hand, the edge context fusion module enhances contextual understanding by merging frequency and spatial features. By integrating complementary information from both domains, the network captures subtle moiré boundaries typically lost in conventional spatial filters. This mechanism

mirrors hybrid models like FSD-Net and MZNet, where multi-domain synergy significantly improved artifact suppression and detail preservation [2][8].

Across evaluation metrics such as Precision, Recall, IoU, and F1-score, the proposed method consistently outperforms established frameworks like MoireDet and MFFNet. At a qualitative level, visual inspections show cleaner, more continuous boundary predictions with minimal false edge detections even across varying light or texture conditions. These outcomes affirm the framework's ability to distinguish real edges from periodic moiré signals—a crucial advantage in practical visual restoration workflows.

6. Local Self-Government Context and Problem Formulation

6.1 From image quality to public-service quality

Information from roads, public buildings, drainage, waste facilities, street assets and other common spaces is used by local authorities to inform their operational decisions. Images can be provided by field inspectors, fixed cameras, mobile devices, contractors or residents. If something is visually inconsistent, it reduces the evidence available to a municipal officer, and vice versa; if it seems to be a physical boundary, it may be just that. If it is inconsistent visually, that reduces the evidence available to a municipal officer, and vice versa, if it looks like a physical boundary, it could be just that. What this means is not that a computer-vision model will directly control a municipality, but that bad image data will skew prioritisation and/or re-inspections and/or avoidable disagreements between a resident and an administration. As such, a reliable boundary map can serve as a quality-control indicator for the human workflow.

6.2 A bounded use case

One viable municipal application might be images submitted for initial infrastructure assessment. An image-quality module can identify possible moiré areas, generate a confidence map, and call for reacquisition of the image or for a human review if the confidence is low. It should not independently consider a road to be unsafe, hold the driver liable, refuse a service request or initiate a sanction. The local authority is responsible for the purpose of collection, evidence threshold, response time and appeal route. This is a bounded role which requires the technical system to remain dependent upon administrative judgment, and have its contribution audited.

6.3 Institutional fit

The workflow involves the team from the service department that owns the case, the information-technology team that maintains the model, the records or data-protection officer, procurement persons and an internal audit or grievance function. The municipality should record who has access to the source images, who can adjust thresholds, how changes to the model are approved and how citizens can challenge an image based decision. These arrangements are similar to the broader literature on smart cities, where data governance is defined as an institutional responsibility, data quality, data security, and public trust [10–14].

7. Proposed Governance-Aware Framework

7.1 Design principles

The framework adds governance controls around, rather than inside, the original neural architecture. First, purpose limitation defines the service problem for which images may be analysed. Second, data minimisation limits collection and retention to what is necessary. Third, uncertainty-aware triage distinguishes confident predictions from cases requiring a second look. Fourth, human accountability ensures that a named municipal officer makes consequential decisions. Fifth, traceability records the model version, threshold, reviewer action, and final service outcome. Finally, accessibility and contestability provide a route for residents to correct records or request reconsideration.




Figure 9. Governance-aware municipal visual-data workflow (proposed; not an evaluated deployment).

7.2 Technical-to-administrative interface

The existing multi-source triplet training and edge-context fusion approach produces a predicted moiré edge map. For a municipal deployment, that output should be accompanied by a calibrated confidence estimate, an image-quality status, and a concise reason code. A low-confidence result should lead to manual review or recapture rather than an automatic adverse action. An accepted image can proceed to the relevant service workflow, with the original image and transformed output retained only under a documented retention schedule. The proposed interface is a design recommendation; confidence calibration, integration costs, and staff workload have not been measured in the current experiments.




Figure 10. Municipal accountability lifecycle for automated image evidence (proposed).

7.3 Workflow and interpretation

Figure 9 presents the proposed governance-aware workflow. The technical model is positioned between data-quality checks and human verification. This placement prevents a common implementation error: treating a high-performing perception component as if it were a complete administrative decision system. The authority must also define an escalation path, record overrides, monitor false alarms, and periodically review whether image collection remains proportionate to the stated public purpose.

8. Governance Evaluation Framework and Indicators

8.1 Evaluation beyond pixel metrics

Precision, recall, IoU, and F1-score remain appropriate for the boundary-detection task, but they cannot establish whether a municipal workflow is fair, transparent, lawful, or useful. A local-government pilot should therefore report two separate scorecards: a technical scorecard for image analysis and a governance scorecard for institutional performance. Combining them into a single headline score would conceal trade-offs and imply comparability where none has been established.

8.2 Proposed indicators

Table 4 provides candidate indicators for a future field evaluation. The indicators are proposed measures, not findings from the present dataset. Baselines and targets should be agreed with the participating local authority before deployment, and disaggregated reporting should be used only where lawful and privacy-preserving.

Table 4. Proposed governance and service indicators for a future municipal pilot; these indicators have not been measured.

Dimension	Candidate indicator	How to interpret
Technical quality	Precision, recall, IoU, F1-score	Report per dataset and relevant scene/device subgroup.
Image handling	Share of flagged images; recapture rate	Check whether flags help identify unusable evidence.
Administrative efficiency	Median time to triage; staff minutes per case	Compare with a baseline workflow; include review overhead.
Accountability	Share of cases with recorded human review	Audit completeness and override reasons.
Contestability	Time to resolve image-related complaints	Check access to correction and alternative channels.
Equity	Variation in rejection/recapture rates	Investigate disparities by context without unnecessary profiling.
Privacy	Retention compliance; unauthorised access incidents	Track access logs, deletion and incident response.
Resilience	Availability; manual fallback readiness	Ensure service continuity during model/system failure.

8.3 Decision rule

A deployment decision should require both acceptable technical performance and acceptable governance performance. For example, a system with strong edge IoU but poor correction procedures should not be considered ready for public-service use. Similarly, a well-documented process cannot compensate for boundary predictions that are too unreliable for the intended task. A staged pilot, with an initial shadow mode in which predictions do not affect service decisions, can reveal failure patterns before operational use.

9. Risks, Ethics, and Accountability in Municipal Deployment

Table 5. Suggested responsibility allocation for accountable municipal deployment.

Actor	Primary responsibility	Minimum record or control
Service department	Define service purpose and decision thresholds	Case protocol; review and escalation route
IT/data team	Maintain model, access and logging	Version record; security and drift checks
Data-protection/records role	Assess lawful basis, minimisation and retention	Impact assessment; retention schedule
Frontline reviewer	Validate uncertain outputs and correct errors	Review outcome and override reason
Internal audit/grievance function	Review complaints and procedural fairness	Periodic audit; complaint-resolution record
Vendor/contractor, if used	Meet contracted security and documentation duties	Change log; incident notice; exit/export plan

9.1 Privacy and function creep

Images captured for infrastructure maintenance may incidentally contain faces, vehicle registration plates, home entrances, or sensitive activities. Local authorities should define whether such content is needed, apply masking where feasible, restrict access, and prohibit unrelated secondary use unless a separate legal and public-interest assessment authorises it. A technically possible use is not automatically a legitimate municipal purpose. The data-governance literature highlights the importance of clear responsibilities, access rules, trust, and protection against function creep [11–14].

9.2 Unequal error burdens

Image quality may vary by neighbourhood, device, weather, lighting, or the resources available to residents. If residents in poorly connected areas use older devices, a strict image-quality threshold could create more recapture requests or delay service. Authorities should compare rejection and recapture rates across relevant service contexts, provide an alternative reporting channel, and avoid making access to a service conditional on the ability to submit a technically ideal image.

9.3 Explainability and contestability

A citizen-facing explanation need not expose model internals. It should state that the image was flagged for possible interference, explain what the flag means for the case, and offer practical next steps. Internal records should preserve the source image where lawful, model version, confidence score, reviewer identity or role, and the reason for any override. An accessible complaint process is essential where automated image analysis contributes to a delay or disputed decision.

9.4 Procurement and vendor accountability

If a municipality purchases the model as a service, the contract should address data ownership, permitted reuse, retention, incident reporting, security testing, model updates, audit access, and exit arrangements. Procurement criteria should include documentation and service support, not only benchmark accuracy. The authority should retain the ability to suspend the model if performance drifts or if the vendor cannot explain a material update.

9.5 Environmental and operational resilience

Municipal systems should record processing time, hardware requirements, availability, and fallback arrangements. A service should continue through a manual route during outages. Any claim that the model reduces cost or accelerates service delivery should be tested against staff time, recapture rates, maintenance overhead, and the cost of errors; those impacts are not inferred from the present computer-vision scores.

10. Relationship to Related Image-Analysis Research

10.1 Transferable lessons, not direct validation

Prior work by Selvaraju and Rajamani on underwater image enhancement demonstrates a restoration-oriented use of modified residual blocks within an improved CycleGAN generator [25]. Karthik, Annapoorani, and Dineshkumar address moving-object recognition and tracking in underwater sonar images [26]. These studies are relevant as examples of image enhancement and visual tracking under challenging acquisition conditions. They do not evaluate moiré boundary detection or municipal governance, and they should not be presented as direct evidence for the proposed local-government application.

10.2 Why the distinction matters

The underwater papers support the broader methodological point that image acquisition conditions can affect downstream visual interpretation. Their settings, sensors, targets, and evaluation objectives differ from municipal imagery, however. Any transfer to local-government services would require a new dataset, task-specific annotation, validation across devices and neighbourhood conditions, and review of public-sector legal and ethical requirements. This separation of technical precedent from governance evidence helps avoid overstating what prior publications establish.

11. Implementation Roadmap for a Local Authority

11.1 Stage one: service definition

Select one low-risk service process and document the problem, legal basis, responsible department, data flow, and expected public benefit. Consult frontline staff and, where appropriate, residents before collecting or reusing imagery.

11.2 Stage two: shadow evaluation

Run the detector without allowing outputs to determine case outcomes. Compare predictions with independently annotated boundaries, record uncertainty, and examine errors by image source, lighting, scene type, and device. Measure whether flags actually help staff identify images that need recapture.

11.3 Stage three: controlled pilot

If the shadow evaluation is satisfactory, introduce the output as an advisory signal with mandatory human review for consequential cases. Publish a short notice explaining the purpose and limits of the tool, and provide a non-digital alternative for submitting service requests.

11.4 Stage four: periodic review

Review performance, complaints, overrides, false alarms, staff workload, and retention compliance at predefined intervals. Pause or revise the system if error rates rise, if disparities emerge, or if the original purpose changes. A public report should distinguish model-level performance from administrative outcomes and disclose limitations in accessible language.

12. Limitations and Future Research

12.1 Current evidence boundary

The manuscript's reported numerical results concern image-boundary metrics on MoireScape and MoireIDT, as stated in the existing sections. No municipal case study, staff evaluation, citizen survey, cost analysis, privacy impact assessment, or public-service outcome dataset is supplied in the source manuscript. The governance framework in Sections 6–11 is consequently a proposed application model and not an empirically validated local-government intervention.

12.2 Future research design

A next study should partner with one or more local authorities and preregister the service use case, inclusion criteria, ground-truth protocol, technical metrics, governance indicators, and stop conditions. The evaluation should compare the normal workflow with a human-led workflow assisted by image-quality flags. Suitable outcomes include repeat-image requests, time to triage, staff review burden, error correction, complaint resolution, and differences across image sources. Interviews with staff and residents can help determine whether the workflow is understandable and trusted. A comparative design across municipalities would help distinguish model effects from differences in staffing, procurement, and administrative practice.

12.3 Contribution to local self-government

The main contribution to the journal's scope is not a claim that moiré detection itself changes local governance. It is the argument that a narrow computer-vision component should be embedded in an accountable public-service arrangement, with explicit responsibility, human review, transparent limits, and avenues for correction. This framing makes the technical system a subject of local administrative design and critical assessment rather than a substitute for public authority.

13. Policy Implications for Local Self-Government

The proposed framework has three implications for municipal policy. First, local authorities should treat automated image analysis as a supporting administrative capability rather than as an autonomous decision-maker. The distinction matters because local self-government involves more than operational efficiency: public decisions must remain attributable to an institution, open to explanation, and capable of correction. Procurement specifications and internal procedures should state explicitly which decisions may use model outputs, which require human confirmation, and which must never be automated on the basis of a boundary map alone.

13.1 Capacity and interdepartmental coordination

A municipality does not need to build a large artificial-intelligence unit to pilot an image-quality component, but it does need named responsibility and sufficient staff capacity. Service departments understand the operational problem; technical staff understand data pipelines and model limitations; records and privacy personnel understand retention and access; and grievance officers understand how errors affect residents. A small cross-functional group can approve the use case, review the risk register, define acceptance criteria, and decide when the system must be paused. This is particularly important where local authorities rely on external vendors or shared regional platforms.

13.2 Transparency without overclaiming

Public communication should describe the system's narrow purpose, the types of images processed, the role of human review, the retention period, and the route for correcting or challenging a record. The authority should not describe the system as "objective" simply because it is algorithmic. Performance varies with data quality and context, and the administrative meaning of a prediction depends on local procedures. Published dashboards should therefore separate model measures from service measures and include limitations, rather than presenting a single aggregate score as proof of successful governance.

13.3 A proportionate adoption principle

Adoption should proceed only where the expected public benefit is specific and testable, less intrusive alternatives have been considered, and the authority can maintain a manual fallback. For a low-risk image triage task, an advisory flag may be proportionate; using the same output to determine enforcement or

eligibility would require a much stronger legal, technical, and democratic justification. The framework thus supports a graduated approach: define the service need, evaluate in shadow mode, conduct a limited pilot, publish lessons, and expand only after the evidence supports expansion. This approach connects innovation to administrative responsibility and helps local institutions retain control over the technologies they deploy.

14. Conclusion and Future Work

14.1 Conclusion

This work presented an innovative framework for accurate moiré boundary detection using multi-source triplet data combined with an edge context fusion mechanism. By leveraging aligned triplet datasets—comprising natural clean images, isolated moiré layers, and synthetic mixture images—and fusing spatial and frequency domain features within a multi-encoder architecture, the proposed model achieves high precision in delineating moiré boundaries. Extensive evaluation on public datasets such as MoireScape and MoireIDT demonstrated significant improvements in boundary detection metrics (Precision, Recall, F1-score, and IoU) when compared to state-of-the-art methods including MoireDet and MFFNet. The triplet training strategy ensured robust pixel-level discrimination by contrasting moiré features across multiple data perspectives, while the edge fusion approach effectively captured subtle variations in moiré patterns by combining complementary signal domains. This joint mechanism not only enhanced detection accuracy but also strengthened model generalization to diverse imaging conditions and device types, including smartphone cameras and printed materials.

Furthermore, the model maintains a balanced computational efficiency, making it viable for integration into consumer-facing imaging devices or real-time demoiréing pipelines. The qualitative improvements in visual moiré edge maps also open pathways for more targeted artifact suppression and higher-fidelity image restoration.

14.2 Future Work

Future work should explore self-supervised learning frameworks to alleviate dependence on costly, well-aligned triplet datasets by exploiting inherent image statistics and corruption patterns for moiré feature learning, as suggested by recent methods like SIDME [8]. Additionally, integrating moiré boundary detection within end-to-end demoiréing pipelines via multi-task learning could jointly optimize artifact detection and removal, improving restoration fidelity [1]. To enable real-time applications, investigations into lightweight architectures through pruning, quantization, or neural architecture search can produce efficient models deployable on mobile and embedded platforms [2]. Moreover, extending the framework for temporal consistency in video sequences may enhance moiré suppression in dynamic content [3]. Expanding datasets with diverse moiré scenarios, including high-frequency textiles and medical imaging, alongside synthetic data generation via adversarial approaches, will enhance model robustness and applicability in practical domains.

References

1. Yang, C., Yang, Z., Ke, Y., Chen, T., Grzegorzec, M., & See, J. (2023). Doing More with Moiré Pattern Detection in Digital Photos. *IEEE Transactions on Image Processing*, 32, 694–708.
2. Lee, S., Baek, W., Kim, Y., Moon, H., Yoo, D., & Park, E. (2025). Moiré Zero: An Efficient and High-Performance Neural Architecture for Moiré Removal. *arXiv preprint arXiv:2507.22407*.
3. He, M., Chen, Y., & Liu, J. (2019). MopNet: Moiré Pattern Detection Using Multi-scale Feature Fusion. *IEEE ICCV*.
4. Liu, H., et al. (2024). Edge-Aware Multi-Feature Fusion for Robust Moiré Boundary Detection. *Proceedings of CVPR*.
5. Photography Life. (2022). What is Moiré and How it Can Ruin Your Photos.
6. ProEdu. (2024). Understanding Moiré in Photography and How to Avoid It.
7. Skylum. (2024). What Is A Moiré Pattern And Why It Can Ruin Your Photo.

8. Wang, F., Zhao, H., & Song, J. (2025). MFFNet: Joint Demoiréing and Super-Resolution. ScienceDirect.
9. Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A Unified Embedding for Face Recognition and Clustering. CVPR.
10. Meijer, A., & Rodríguez Bolívar, M. P. (2016). Governing the smart city: A review of the literature on smart urban governance. *International Review of Administrative Sciences*, 82(2), 392–408. <https://doi.org/10.1177/0020852314564308>
11. OECD. (2023). *Smart City Data Governance: Challenges and the Way Forward*. OECD Publishing. <https://doi.org/10.1787/e57ce301-en>
12. Choenni, S., Bargh, M. S., Busker, T., & Netten, N. (2022). Data governance in smart cities: Challenges and solution directions. *Smart Cities and Society*, 1(1). <https://doi.org/10.3233/SCS-210119>
13. van Zoonen, L. (2021). Citizen-centered data governance in the smart city: From ethics to accountability. *Sustainable Cities and Society*, 75, 103308. <https://doi.org/10.1016/j.scs.2021.103308>
14. Landsbergen, D., Girth, A., & Westover-Muñoz, A. (2022). Governance rules for managing smart city information. *Urban Governance*. <https://doi.org/10.1016/j.ugj.2022.05.003>
15. Kitchin, R., Lauriault, T. P., & McArdle, G. (2015). Knowing and governing cities through urban indicators, city benchmarking and real-time dashboards. *Regional Studies, Regional Science*, 2(1), 6–28. <https://doi.org/10.1080/21681376.2014.983149>
16. Kitchin, R. (2014). The real-time city? Big data and smart urbanism. *GeoJournal*, 79, 1–14. <https://doi.org/10.1007/s10708-013-9516-8>
17. Hollands, R. G. (2008). Will the real smart city please stand up? *City*, 12(3), 303–320. <https://doi.org/10.1080/13604810802479126>
18. Nam, T., & Pardo, T. A. (2011). Conceptualizing smart city with dimensions of technology, people, and institutions. *Proceedings of the 12th Annual International Digital Government Research Conference*, 282–291. <https://doi.org/10.1145/2037556.2037602>
19. Gil-Garcia, J. R., Zhang, J., & Puron-Cid, G. (2016). Conceptualizing smartness in government: An integrative and multi-dimensional view. *Government Information Quarterly*, 33(3), 524–534. <https://doi.org/10.1016/j.giq.2016.03.002>
20. Cardullo, P., & Kitchin, R. (2019). Being a “citizen” in the smart city: Up and down the scaffold of smart citizen participation in Dublin, Ireland. *GeoJournal*, 84, 1–13. <https://doi.org/10.1007/s10708-018-9845-8>
21. Batty, M., Axhausen, K. W., Giannotti, F., et al. (2012). Smart cities of the future. *The European Physical Journal Special Topics*, 214, 481–518. <https://doi.org/10.1140/epjst/e2012-01703-3>
22. Janssen, M., & Estevez, E. (2013). Lean government and platform-based governance—Doing more with less. *Government Information Quarterly*, 30, S1–S8. <https://doi.org/10.1016/j.giq.2012.11.003>
23. Zuiderwijk, A., Janssen, M., & Davis, C. (2014). Innovation with open data: Essential elements of open data ecosystems. *Information Polity*, 19(1–2), 17–33. <https://doi.org/10.3233/IP-140329>
24. Marchesani, F., & Ceci, F. (2026). Rules and responsibilities in smart city governance: Exploring the architecture of data value chain. *Government Information Quarterly*, 43(1), 102115. <https://doi.org/10.1016/j.giq.2026.102115>
25. Selvaraju, K., & Rajamani, S. (2024). New approach to underwater image enhancement using modified residual blocks in generator architecture for improved cycle generative adversarial networks. *Proceedings of the Bulgarian Academy of Sciences*, 77(1), 73–81. <https://doi.org/10.7546/CRABS.2024.01.09>
26. Karthik, S., Annapoorani, V., & Dineshkumar, S. (2016). Recognition and tracking of moving object in underwater sonar images. *International Journal of MC Square Scientific Research*, 8(1), 93–98. <https://doi.org/10.20894/IJMSR.117.008.001.010>