



ARTICLE

High-Fidelity Characterization and Physical Quantification of Bridge Pier Cracks Based on Oriented Object Detection

Zongqi Xiong¹, Yu Li² and Xinxin Cao^{3,*}

¹JSTI Group Chongqing Inspection, Testing and Certification Co., Ltd., Chongqing, China

²Chongqing Highway Maintenance Engineering (Group) Co., Ltd., Chongqing, China

³Beijing University of Technology, Beijing, China

*Corresponding Author: Xinxin Cao. Email: caoxinxin@sdut.edu.cn

Received: 08 April 2026; Accepted: 20 May 2026; Published: 23 October 2026

ABSTRACT: Accurate detection and quantitative characterization of bridge pier cracks are essential for UAV-assisted structural health monitoring, but conventional horizontal bounding box detectors often suffer from geometric mismatch, background redundancy, and missed detection of slender cracks in dense regions. To address these limitations, this study proposes an oriented object detection-based framework for high-fidelity crack representation and physical quantification. A dedicated PierCrack-OBB dataset was first constructed from real bridge inspection images and annotated using multi-segment oriented bounding boxes to better represent locally curved and high-aspect-ratio cracks. On this basis, the standard YOLOv8-OBB framework was improved by incorporating DCNv3 for adaptive crack feature extraction, a channel-spatial attention module for background suppression, and Gaussian probability loss for stable rotated-box regression. In addition, the minor-axis and inclination parameters of the detected oriented boxes were combined with the ground sample distance of UAV images to estimate physical crack width and describe crack orientation characteristics. Experimental results show that the proposed method increased mAP@0.5 from 88.6% for standard YOLOv8-OBB to 91.6%, while achieving 92.8% precision and 90.1% recall. For crack width estimation, the mean absolute error remained below 0.18 mm. These results demonstrate that the proposed method improves both the geometric fidelity of crack detection and the engineering interpretability of detection outputs. The main novelty of this study lies in integrating multi-segment OBB-based crack representation, improved oriented detection, and GSD-based physical quantification into a unified framework, thereby extending crack detection from image-level localization to measurable width estimation and orientation characterization.

KEYWORDS: Bridge pier; crack detection; oriented object detection; geometric characterization; physical quantification; structural health monitoring

1 Introduction

Bridges are critical components of transportation infrastructure, and their service safety and operational condition are directly related to traffic efficiency and public safety [1]. Among bridge substructure components, piers transmit superstructure loads and maintain structural stability; therefore, pier deterioration may significantly reduce the overall safety reserve of a bridge [2]. With increasing service time, concrete bridge piers are prone to different types of surface cracking under the combined effects of mechanical loading, environmental exposure, material aging, and fatigue [3]. Among these defects, vertical cracks are commonly regarded as important visual indicators for assessing the compression condition, local damage evolution, and durability degradation of bridge piers [4]. Therefore, timely, accurate, and quantitative

identification of pier cracks is of considerable importance for condition assessment, damage warning, and maintenance decision-making in bridge structural health monitoring [5]. Conventional bridge inspection still relies heavily on close-range manual visual examination, which is labor-intensive, time-consuming, and susceptible to environmental conditions and inspector experience [6]. In recent years, with the rapid development of unmanned aerial vehicle (UAV) platforms, high-resolution imaging devices, and intelligent sensing technologies, machine vision-based non-contact inspection methods have attracted increasing attention and provided a new technical route for efficient and safe bridge inspection [7].

Although deep learning has achieved rapid progress in crack detection tasks, most existing object detection approaches still rely on horizontal bounding boxes (HBB) for damage localization [8]. Such methods perform well in general object detection scenarios, but they exhibit limited geometric representation capability when applied to bridge pier cracks, which are typically slender, curved, irregularly oriented, and locally dense [9]. On the one hand, HBBs cannot accurately fit the actual contour of cracks and often include a large amount of irrelevant background, thereby weakening the model's ability to learn crack-specific features [10]. On the other hand, in regions where multiple cracks are closely distributed or intersected, horizontal boxes tend to overlap substantially, which may lead to missed detections or false detections during non-maximum suppression [11]. As a result, HBB-based representation is constrained by both geometric mismatch and background interference in bridge pier crack detection, limiting the robustness and reliability of detection results in complex engineering environments [12].

Beyond detection accuracy, HBB-based methods are also limited in describing crack geometry. For bridge pier cracks, geometric attributes such as propagation direction, spatial distribution, and local width not only characterize their surface morphology but may also provide supportive information for interpreting structural stress conditions and damage patterns [13]. For instance, vertical cracks are often associated with the development of transverse tensile strain in compression members, while diagonal cracks may be related to shear effects or localized stress concentration [14]. However, HBBs only provide positional information and axis-aligned dimensions, making it difficult to describe crack inclination and local directional variations. Consequently, they are not well suited to the demand for high-fidelity geometric characterization and subsequent quantitative analysis in bridge engineering applications.

Oriented object detection based on oriented bounding boxes (OBB) offers a promising solution to the above limitations. Compared with HBBs, OBBs introduce an additional angle parameter, enabling a more flexible fit to the actual boundary of slender targets in the image plane, thereby reducing background redundancy and improving target separation in dense scenarios [15]. Although segmentation-based methods can provide more detailed pixel-level crack contours, they usually require labor-intensive pixel-wise annotations and are more sensitive to low contrast, weathering texture, shadows, and ambiguous crack boundaries in UAV inspection images. In this study, OBB detection was selected because it provides a practical balance between annotation cost, detection efficiency, geometric fidelity, and engineering quantification. The multi-segment OBB representation can preserve the local direction and effective width of slender cracks while directly providing major-axis, minor-axis, and inclination-angle parameters for subsequent physical width estimation and orientation analysis. This characteristic makes OBBs particularly suitable for bridge pier crack detection. In addition, the geometric parameters contained in OBBs, including the major axis, minor axis, and inclination angle, provide useful information for crack width estimation, orientation statistics, and damage feature analysis. However, studies specifically focusing on bridge pier crack detection using OBBs remain limited. In particular, systematic research is still lacking in terms of dedicated OBB dataset construction, feature extraction for slender and non-rigid cracks, and the conversion of detection results into engineering-related physical quantities.

Nevertheless, most existing studies still treat crack detection and crack quantification as two relatively independent tasks. In many vision-based inspection frameworks, object detectors are mainly used to determine whether and where cracks exist, while the subsequent engineering interpretation of crack width, orientation, and damage-related geometric features is usually performed through separate post-processing, segmentation, or manual measurement procedures. As a result, the geometric information contained in detection outputs is not fully exploited for physical-scale crack assessment. This limitation is particularly evident for bridge pier cracks, because these cracks are often slender, locally curved, densely distributed, and irregularly oriented. Although oriented bounding boxes can provide angle-aware and scale-aware geometric parameters, their potential for simultaneously improving detection fidelity and supporting physical quantification of crack width and orientation has not been systematically investigated in bridge pier inspection scenarios. Therefore, the core research gap addressed in this study is the lack of an integrated framework that connects high-fidelity oriented crack detection with engineering-oriented physical quantification under practical UAV inspection conditions.

To bridge this research gap, this study proposes a high-fidelity geometric characterization and physical quantification method for bridge pier cracks based on oriented object detection. First, a dedicated PierCrack-OBB dataset was established using images collected from real bridge service environments, and cracks were annotated with multi-segment oriented bounding boxes. Second, based on the YOLOv8-OBB framework, deformable convolution, a channel-spatial attention mechanism, and Gaussian probability loss were introduced to improve feature extraction and oriented regression for slender, curved, and non-rigid crack targets. Finally, by combining the geometric parameters of OBBs with the ground sample distance (GSD) of UAV inspection images, a mapping was established from OBB-based detection results to physical crack width and orientation-related characteristics, enabling further quantitative expression of crack features. The present study aims to provide a feasible technical approach for automatic crack identification, geometric characterization, and auxiliary assessment of bridge piers, and to support intelligent inspection research in bridge structural health monitoring.

To clarify the novelty and technical contributions of this study, the main contributions are summarized as follows. First, a dedicated PierCrack-OBB dataset was constructed from real UAV-based bridge pier inspection images, and multi-segment OBB annotations were adopted to represent slender, locally curved, and irregularly oriented cracks with higher geometric fidelity than conventional HBB annotations. Second, an improved YOLOv8-OBB framework was developed for high-aspect-ratio and non-rigid crack targets by integrating DCNv3, a channel-spatial attention module, and Gaussian probability loss, thereby improving adaptive crack feature extraction, background suppression, and rotated-box regression stability. Third, an engineering-oriented physical quantification strategy was established by combining OBB geometric parameters with GSD, enabling crack width estimation and orientation characterization directly from detection outputs. Therefore, the novelty of this work lies not only in improving oriented crack detection accuracy, but also in connecting high-fidelity OBB-based localization with physical-scale crack quantification for bridge pier inspection.

2 Methods

2.1 Construction of the PierCrack-OBB Dataset for Engineering Applications

In bridge structural health monitoring, data quality directly affects model performance and its applicability to real engineering scenarios. To improve the task specificity and generalization capability of crack detection for bridge substructures, a dedicated dataset, namely PierCrack-OBB, was constructed in this study. During data acquisition, an industrial UAV equipped with a high-resolution optical sensor was employed to inspect bridge piers under long-term environmental exposure and operational traffic loading. To make the

dataset source clearer, the PierCrack-OBB images were collected from multiple in-service highway bridges and included multiple bridge piers in total. To improve the representativeness of the PierCrack-OBB dataset, the image samples were collected from multiple in-service bridge pier inspection scenes rather than from a single controlled laboratory environment. The acquisition targets covered common bridge pier forms in practical highway bridge inspection, including rectangular piers with flat surfaces and cylindrical or rounded piers with curved surfaces. The dataset also included pier surfaces with different material and deterioration conditions, such as relatively clean concrete surfaces, weathered concrete, water-stained regions, locally shadowed areas, rough surface textures, and areas with dense or intersecting crack patterns.

In terms of imaging scenes, the collected images covered close-range frontal views, upward-looking views near the pier cap, downward-looking views near the pier base, and oblique views caused by UAV pose variation and field accessibility constraints. The UAV shooting distance was mainly controlled within 2–4 m from the pier surface to balance image resolution, inspection safety, and flight operability. The images were acquired under natural outdoor lighting conditions, including normal illumination, local shadows, and slight illumination variation caused by UAV viewing angle and pier-surface geometry. The crack samples covered vertical cracks, diagonal cracks, horizontal or local flexural cracks, map-like or intersecting cracks, and irregular surface cracks. Among these conditions, approximately frontal images were used to provide stable geometric references, while oblique and perspective-distorted images were retained to improve the robustness of the dataset under realistic UAV inspection conditions. The typical inspection scenario and imaging geometry are shown in Fig. 1, where Fig. 1a presents the close-range UAV inspection process and Fig. 1b illustrates the geometric relationship among the shooting distance, focal length, and image resolution.




Figure 1: UAV-based bridge pier inspection scenario and GSD calculation geometry: (a) close-range UAV inspection; (b) imaging geometry used for GSD calculation.

To satisfy the scale accuracy required for physical crack quantification, a mapping from image pixels to real-world dimensions was established based on the camera imaging model. The ground sample distance (GSD) was calculated as follows:

$$GSD = \frac{D \times 10^3 \times S_w}{f \times N_w} \quad (1)$$

where GSD is the ground sample distance in mm/pixel, D is the UAV shooting distance in m, S_w is the physical width of the camera sensor in mm, f is the effective focal length in mm, and N_w is the horizontal image resolution in pixels. In this study, $S_w = 13.2$ mm, $f = 8.8$ mm, and $N_w = 5472$ pixels. To balance image clarity and operational safety, the shooting distance was controlled within 2–4 m from the pier surface. At

a distance of 3 m, one pixel corresponds to approximately 0.82 mm/pixel, while at 1.8 m, the GSD reaches about 0.5 mm/pixel, which is sufficient for millimeter-level crack quantification. Before width conversion, the UAV images were first corrected using the camera intrinsic parameters to reduce lens distortion. For upward-looking, downward-looking, and oblique views, the possible influence of perspective distortion was considered during width validation. Images with severe perspective distortion or uncertain local surface orientation were excluded from the primary physical width validation samples. Considering that bridge piers usually have cylindrical or rectangular cross-sections, the dataset includes not only frontal views but also upward, downward, and perspective-distorted views, so as to better represent practical inspection conditions. Because the GSD calculation assumes an approximately perpendicular imaging geometry and a calibrated camera, deviations in shooting distance, camera pose, pier-surface curvature, perspective distortion, and residual lens distortion may introduce local scale errors in the pixel-to-millimeter conversion. These factors were therefore considered as uncertainty sources in the subsequent crack-width error analysis. These images realistically capture the surface deterioration of bridge piers under long-term environmental actions and provide engineering-representative samples for subsequent model training [16].

After image collection, high-quality crack annotation is essential for geometric representation and subsequent quantitative analysis. Conventional horizontal bounding boxes are not well suited to the slender, curved, and locally deflected geometry of pier cracks, and they usually introduce substantial background redundancy. To address this issue, a fine-grained annotation strategy based on multi-segment oriented bounding boxes was adopted [17]. Specifically, each continuous crack was divided into several locally linear OBB segments according to changes in crack direction, local curvature, width variation, discontinuity, or branching. In approximately straight regions, a relatively long OBB was used to cover the continuous crack segment, whereas in bending regions, several shorter OBBs were used to better follow the local crack path. The endpoints of each segment were selected along the visible crack centerline, and a new OBB segment was created when the local direction changed obviously or when one box could no longer fit the crack contour without including excessive background.

For each OBB segment, the major axis was aligned with the local crack propagation direction, and the minor axis was aligned with the local width-normal direction. The OBB angle was determined by the direction of the local crack centerline within the segment, while the box width was adjusted to cover the visible transverse crack opening. The fitting error was defined as the maximum visible deviation between the OBB boundary and the crack contour in clearly identifiable edge regions; annotations with a boundary deviation greater than approximately 2 pixels were adjusted by shortening the segment or modifying the OBB angle and width.

Adjacent OBBs belonging to the same crack were allowed to slightly touch or overlap at their ends to maintain crack continuity, but excessive overlap and one-box coverage of multiple neighboring cracks were avoided. To ensure annotation consistency, all images were labeled according to the same annotation protocol. After initial labeling, the annotations were reviewed through a second manual cross-check, focusing on segment continuity, OBB angle consistency, local width coverage, excessive background inclusion, and missing segments in curved or branched regions.

Before formal annotation, all annotators were trained using the same representative crack examples, and inconsistent cases were discussed to unify the segmentation criteria, angle definition, and boundary-fitting rules. During annotation review, inter-annotator variability was evaluated by checking whether different annotators produced consistent crack segmentation, OBB orientation, local width coverage, and boundary deviation for the same or similar crack patterns. Discrepant annotations were corrected through consensus review, and annotations exceeding the approximately 2-pixel boundary deviation criterion were revised before model training.

To make the data source more transparent, the dataset was further summarized from the perspectives of image scale, annotation quantity, aspect-ratio distribution, imaging distance, and scene diversity. The image acquisition covered multiple viewing directions and practical background conditions commonly encountered in UAV-based bridge inspection. Specifically, the dataset retained samples with frontal, upward, downward, and oblique imaging geometries, as well as samples affected by weathering textures, illumination variation, water stains, shadows, and local surface roughness. These scene-level variations were intentionally included to improve the generalization ability of the detector and to better reflect the complexity of real bridge pier inspection in [Table 1](#).

Table 1: Acquisition and scene characteristics of the PierCrack-OBB dataset.


Item	Description
Acquisition platform	Industrial UAV with a high-resolution optical sensor
Inspection object	In-service bridge piers under operational traffic and environmental exposure
Pier geometry	Rectangular piers with flat surfaces; cylindrical or rounded piers with curved surfaces
Imaging distance	Mainly 2–4 m from the pier surface
Imaging viewpoints	Frontal, upward-looking, downward-looking, and oblique views
Surface conditions	Clean concrete, weathered concrete, water stains, local shadows, rough textures, and surface contamination
Crack morphology	Vertical cracks, diagonal cracks, horizontal/local cracks, dense cracks, intersecting cracks, and irregular map-like cracks
Image resolution	5472 pixels in the horizontal direction
Annotation form	Multi-segment oriented bounding boxes
Dataset size	2542 images and 12,245 OBB annotations

The diversity of the PierCrack-OBB dataset was considered from both structural and imaging perspectives. Structurally, the dataset includes common in-service highway bridge pier forms, such as rectangular piers with flat surfaces and cylindrical or rounded piers with curved surfaces. In terms of surface and environmental conditions, the images cover clean concrete, weathered concrete, water-stained regions, local shadows, rough textures, surface contamination, and dense or intersecting crack patterns. From the imaging perspective, the dataset contains frontal, upward-looking, downward-looking, and oblique UAV views under natural outdoor illumination. Therefore, the dataset provides representative samples for evaluating OBB-based crack detection under typical bridge pier inspection conditions. Nevertheless, it cannot cover all possible bridge types, severe weather conditions, or extreme deterioration states; expanding the dataset across more regions and service environments remains an important direction for future work.

The statistical characteristics of the dataset are listed in [Table 2](#). As shown in [Table 2](#), the PierCrack-OBB dataset contains 2542 images and 12,245 oriented annotations in total. Crack targets exhibit pronounced high-aspect-ratio characteristics, with an average aspect ratio of 26.4, and more than 86% of the samples have aspect ratios greater than 10. As illustrated in [Fig. 2](#), the proposed multi-segment OBB representation can better fit the actual morphology of slender and arbitrarily oriented cracks on bridge piers, while reducing the background redundancy and geometric overlap caused by horizontal boxes. This provides more task-relevant supervision information for the training of oriented crack detection models.

Table 2: Statistics on image scale, annotation features, crack aspect ratio distribution, and acquisition characteristics of the PierCrack-OBB dataset.

Dataset Split	Number of Images	Total OBB Annotations	Average Aspect Ratio	Maximum Aspect Ratio	Proportion of AR > 10
Training set	1779	8542	26.4	82.5	86.20%
Validation set	508	2415	25.8	76.3	84.70%
Testing set	255	1288	27.1	79.1	87.50%
Total/Avg.	2542	12,245	26.4	82.5	86.10%

**Figure 2:** Geometric representation of bridge pier cracks: (a) actual vertical cracks on the surface of an in-service bridge pier; (b) comparison between HBB and multi-segment OBB representations.

In addition to the overall annotation quantity and aspect-ratio statistics, the morphological distribution of crack samples was further analyzed to evaluate the representativeness of the PierCrack-OBB dataset. According to the local propagation direction, spatial morphology, and engineering crack pattern, the annotated crack segments were divided into five categories: vertical cracks, diagonal cracks, horizontal or local flexural cracks, map-like or intersecting cracks, and irregular or other surface cracks. The statistical distribution is summarized in Table 3. It can be observed that vertical cracks account for the largest proportion of the dataset, which is consistent with the common occurrence of longitudinal cracking on bridge pier surfaces. Meanwhile, diagonal cracks, horizontal cracks, and map-like or intersecting cracks are also included, indicating that the dataset covers not only dominant vertical crack patterns but also multiple non-vertical and complex crack morphologies. This distribution improves the representativeness of the dataset and provides diverse training samples for evaluating the adaptability of oriented object detection to different crack geometries.

It should be noted that the crack morphology categories were used mainly to describe the visual and geometric characteristics of the samples. For curved or branched cracks, the category was determined according to the dominant local propagation direction and the overall crack pattern after multi-segment OBB

annotation. Therefore, this classification was intended to support dataset description and model evaluation, rather than to provide a deterministic structural damage diagnosis.

Table 3: Statistical distribution of crack morphology categories in the PierCrack-OBB dataset.

Crack Morphology	Description	Number of OBB Annotations	Proportion (%)
Vertical cracks	Cracks mainly extending along the pier height direction	6972	56.9
Diagonal cracks	Cracks with obvious inclined propagation directions	2018	16.5
Horizontal/local flexural cracks	Cracks mainly extending in the horizontal or near-horizontal direction	1126	9.2
Map-like/intersecting cracks	Randomly distributed, branched, or intersecting crack patterns	1420	11.6
Irregular/other surface cracks	Locally discontinuous or ambiguous crack morphologies affected by surface texture or deterioration	709	5.8
Total	—	12,245	100.0

2.2 Improved Oriented Object Detection Algorithm

After dataset construction and high-fidelity annotation, YOLOv8-OBB was adopted as the baseline detector because it provides a favorable balance between detection accuracy and inference efficiency for oriented object detection. The baseline model extracts multi-scale features through the backbone network, fuses them in the Neck module, and outputs oriented bounding boxes containing the target center, scale, and rotation angle in the detection head. However, bridge pier cracks are typically slender, curved, non-rigid, and locally discontinuous, which limits the adaptability of standard convolution to such targets. To improve the representation capability for pier cracks, several task-oriented modifications were introduced into the baseline model, as shown in Fig. 3.




Figure 3: Overall architecture of the improved YOLOv8-OBB framework for bridge pier crack detection.

First, to enhance the model's ability to perceive local bending, slender edges, and non-rigid propagation patterns of cracks, Deformable Convolution v3 (DCNv3) was incorporated into the deep semantic feature extraction stage of the backbone. In implementation, DCNv3 was inserted into the last two backbone stages

corresponding to the P4 and P5 feature levels, where the feature maps contain stronger semantic information for slender crack targets. The original feature-map resolution and channel dimensions were kept unchanged to maintain compatibility with the YOLOv8-OBB neck and detection head. For standard 2D convolution, the response at position p_0 on the output feature map is obtained by weighted summation over a regular sampling grid. After introducing deformable convolution, the regular sampling positions are adjusted by learnable offsets, and the operation can be written as:

$$y(p_0) = \sum_{k=1}^K w_k x(p_0 + p_k + \Delta p_k) \quad (2)$$

where $y(p_0)$ is the output feature at position p_0 , $x(\cdot)$ is the input feature map, w_k is the convolution kernel weight at the k -th sampling point, p_k is the regular grid offset, Δp_k is the learnable deformable offset, and K is the number of sampling points. Since Δp_k may take continuous values, bilinear interpolation is used to obtain the corresponding feature value. This mechanism enables the sampling positions to adapt to the target morphology, thereby improving the network's sensitivity to curved crack paths and local geometric variations [18].

Second, to suppress the interference caused by concrete weathering, water stains, shadows, and surface texture noise, a channel-spatial attention mechanism was introduced into the FPN+PAN structure of the Neck module. Specifically, the attention module was placed after multi-scale feature fusion and before the detection head, so that the fused P3, P4, and P5 features could be reweighted before final classification and rotated-box regression. The channel attention branch used a reduction ratio of 16, and the spatial attention branch used a 7×7 convolution kernel. The module consists of a channel attention branch followed by a spatial attention branch. Channel attention learns the relative importance of different channels through global average pooling and max pooling, thereby enhancing feature responses that are more sensitive to crack textures. Spatial attention further learns the saliency of different spatial locations, encouraging the network to focus on crack regions while suppressing background interference. Combined with deformable convolution, this attention mechanism helps improve the discrimination of slender cracks in complex weathered backgrounds.

In addition, because bridge pier cracks usually exhibit extreme aspect ratios, conventional IoU-based regression loss for rotated boxes tends to become unstable under small angular perturbations. To improve the angle regression quality for slender targets, a probabilistic loss based on two-dimensional Gaussian modeling of oriented boxes was adopted. The Gaussian probability loss was used as the rotated-box regression term, while the original classification loss and distribution focal loss were retained. Therefore, the total training loss was formulated as $L_{\text{total}} = \lambda_{\text{box}} L_{\text{GPL}} + \lambda_{\text{cls}} L_{\text{cls}} + \lambda_{\text{dfl}} L_{\text{DFL}}$, where L_{GPL} denotes the Gaussian probability loss for oriented-box regression. Specifically, the predicted box and the ground-truth box are represented as two Gaussian distributions, $\mathcal{N}_p(\mu_p, \Sigma_p)$ and $\mathcal{N}_g(\mu_g, \Sigma_g)$, respectively, where the center coordinates are encoded by the mean vector μ , and the width, height, and orientation information are encoded in the covariance matrix Σ . The regression error is measured by the squared 2-Wasserstein distance:

$$D_W^2(\mathcal{N}_p, \mathcal{N}_g) = \|\mu_p - \mu_g\|_2^2 + \text{Tr} \left(\Sigma_p + \Sigma_g - 2 \left(\Sigma_p \left(\frac{1}{2} \Sigma_g^{-1} \right)^{\frac{1}{2}} \right) \right) \quad (3)$$

where $D_W^2(\mathcal{N}_p, \mathcal{N}_g)$ denotes the squared 2-Wasserstein distance between the predicted and ground-truth Gaussian distributions, μ_p and μ_g are the mean vectors corresponding to the predicted and ground-truth box centers, Σ_p and Σ_g are the corresponding covariance matrices, and $\text{Tr}(\cdot)$ denotes the trace operation.

Compared with conventional IoU loss, this probabilistic formulation provides a smoother distance metric for slender rotated objects under small angular variations, thereby improving the stability and robustness of oriented box regression. As shown in Fig. 4, this loss formulation helps alleviate the gradient discontinuity problem when dealing with extremely high-aspect-ratio targets.




Figure 4: Comparison of gradient continuity between traditional IoU loss and Gaussian probability loss under extreme aspect-ratio conditions.

Through the above modifications, the proposed method enhances the representation of slender pier cracks from three aspects: feature extraction, feature selection, and oriented box regression. Specifically, DCNv3 improves the adaptability to non-rigid crack morphology, the channel-spatial attention mechanism strengthens the response to crack regions while suppressing background noise, and the Gaussian probability loss stabilizes the regression of box angle and scale. Together, these modules form an improved oriented object detection framework for bridge pier crack detection.

2.3 Quantification of Crack Geometry and Physical Features

In the proposed framework, the oriented bounding boxes predicted by the detector are not only used for crack localization, but also serve as the basis for quantifying crack geometry and engineering-related physical indicators. To improve the interpretability of the detection results in bridge condition assessment, a mapping from OBB-based pixel parameters to physical crack width and orientation characteristics was established.

First, for crack width quantification, the GSD defined in Section 2.1 was used to convert the minor-axis length of each oriented bounding box from pixels to physical dimensions. Since the annotation protocol specifies that the major axis follows the crack direction and the minor axis is aligned with the local width-normal direction, the minor axis of the detected OBB can be used to approximate the local maximum transverse opening width of the crack segment. This approximation is most reliable for single, locally straight, and well-contrasted crack segments. For branched or obviously curved cracks, the crack path was divided into multiple local OBB segments, and the minor axis was interpreted only as the effective width of the corresponding local segment rather than the width of the entire crack. For irregular, rough-edged, or partially shadowed cracks, the OBB minor axis may include boundary uncertainty or small background regions; therefore, it was used as an engineering-oriented width indicator rather than a strict pixel-level contour

width. By multiplying the minor-axis pixel length by the GSD value of the image, the crack width can be automatically estimated, thus providing a basis for millimeter-level quantitative evaluation of bridge pier cracks.

Second, the inclination angle output by the OBB detector was further used for crack orientation statistics. The spatial direction of cracks on bridge piers may be associated with possible stress states or damage evolution patterns, but such association is not unique or deterministic. Therefore, the angle information in this study was used mainly for geometric orientation description and preliminary auxiliary interpretation, rather than for direct mechanical diagnosis. Based on commonly observed crack orientation patterns in concrete bridge components, empirical angular intervals were defined, as listed in Table 4. These intervals are intended to provide descriptive orientation categories and possible mechanical implications. They should be interpreted together with structural configuration, loading history, material condition, environmental exposure, and on-site inspection evidence, rather than being used as independent diagnostic criteria.

Table 4: Orientation-based crack categories and possible mechanical implications inferred from OBB inclination characteristics.

Orientation Category	Inclination Angle Threshold	Mechanistic Interpretation for Screening	Engineering-Screening Implication
Vertical crack pattern	80°–100°	Compression-related splitting tendency or transverse tensile strain in the pier-height direction	Prioritize bearing-capacity review; check axial compression, eccentric loading, uneven settlement, and local deterioration through field inspection and structural assessment
Diagonal crack pattern	30°–60°/120°–150°	Shear-related stress concentration or combined shear–normal stress state	Prioritize shear-related review; check lateral actions, water-flow impact, seismic effects, local torsion, and diagonal crack continuity
Horizontal/local flexural crack pattern	0°–15°/165°–180°	Local flexural response or tensile stress concentration near edge regions	Prioritize serviceability review; check local bending effects, crack continuity, and reinforcement corrosion risk through follow-up inspection
Map-like or surface crack pattern	Random and intersecting	Shrinkage, temperature effect, weathering, or superficial deterioration pattern	Prioritize durability review; evaluate surface deterioration, crack width, crack depth, moisture ingress, and exposure condition before structural interpretation

As illustrated in Fig. 5, the proposed quantitative mapping model takes the minor-axis length and inclination angle of oriented bounding boxes as the core parameters, enabling the conversion from pixel-level detection results to physical crack width estimation and geometric orientation analysis. In this way, the output of oriented object detection is extended beyond two-dimensional localization and provides richer geometric information for quantitative crack assessment and subsequent engineering analysis of bridge piers.




Figure 5: Quantitative mapping model from oriented bounding box pixel parameters to physical crack width estimation and geometric orientation analysis.

3 Model Performance and Visual Validation

The proposed model was trained and tested under the following hardware and software environments: Intel Xeon W-2255 CPU, NVIDIA RTX A6000 48 GB GPU, and 64 GB RAM. The software platform consisted of Ubuntu 20.04, Python 3.8, PyTorch 1.13, and CUDA 11.7. To improve reproducibility, the random seed was fixed to 42 for Python, NumPy, and PyTorch. The standard YOLOv8 and YOLOv8-OBB models were initialized using official pretrained weights, and the detection heads were adapted to the single crack category in the PierCrack-OBB dataset. For the proposed model, the same YOLOv8-OBB pretrained weights were used for initialization, while the newly introduced modules were trained together with the whole network.

During training, all input images were resized to 640×640 pixels. The batch size was set to 16, and the AdamW optimizer was adopted with a weight decay coefficient of 0.0005. The initial learning rate was set to 0.001 and decayed using a cosine annealing strategy. The model was trained for 300 epochs. The loss weights followed the commonly used YOLOv8-OBB setting, with $\lambda_{\text{box}} = 7.5$, $\lambda_{\text{cls}} = 0.5$, and $\lambda_{\text{dfl}} = 1.5$. These weights were kept the same for all ablation experiments, and the final hyperparameter settings were determined according to validation-set performance without using the independent test set. Data augmentation was applied only to the training set, including random horizontal flipping with a probability of 0.5, random vertical flipping with a probability of 0.2, random rotation within $\pm 15^\circ$, brightness perturbation with a factor of 0.8–1.2, and Gaussian blur with a probability of 0.2. During inference, the confidence threshold was set to 0.25, and the NMS IoU threshold was set to 0.45. Rotated NMS was used for OBB-based models, while standard NMS was used for the HBB baseline. The training, validation, and testing sets strictly followed the split shown in Table 2. Only the training and validation sets were used during model development, while the test set was reserved exclusively for final performance evaluation. To reduce possible data leakage, the dataset

split was performed at the bridge/site, pier-surface, and inspection-area level rather than at the individual-image or crack-instance level. Images collected from the same bridge site, the same pier surface, or highly overlapping UAV inspection regions were assigned to the same subset. Therefore, images from the same local pier area were not simultaneously included in the training and testing sets. The final split ratio was approximately 70% for training, 20% for validation, and 10% for testing, as summarized in [Table 2](#).

3.1 Quantitative Evaluation of Core Detection Performance

To objectively evaluate the performance of the proposed improved oriented object detection model for bridge pier crack detection, Precision, Recall, mean Average Precision at IoU 0.5 (mAP@0.5), and F1-Score were adopted as evaluation metrics. In bridge inspection scenarios, Precision reflects the model's ability to suppress interference from weathering, water stains, and texture noise, whereas Recall indicates its capability to detect actual crack targets. To validate the effectiveness of the proposed method, comparisons were conducted with the conventional horizontal bounding box (HBB)-based baseline model YOLOv8, as well as two mainstream oriented object detection methods, namely RoI Transformer and Oriented R-CNN. The standard YOLOv8-OBB was also included for comparison. To ensure a fair comparison, all compared models were trained and evaluated using the same training, validation, and testing splits. The independent test set was not used during model training or parameter selection. All models used the same input size of 640×640 pixels and were evaluated on the same test set using the same metrics. For the HBB-based YOLOv8 baseline, the original OBB annotations were converted into the minimum enclosing horizontal bounding boxes, while the OBB-based methods directly used the original oriented annotations. During inference, the confidence threshold was set to 0.25 for all models, and the NMS IoU threshold was set to 0.45. Standard NMS was used for the HBB baseline, whereas rotated NMS was used for OBB-based detectors. The quantitative results are listed in [Table 5](#).

Table 5: Quantitative comparison of detection accuracy on the PierCrack-OBB test set.

Detection Model	Bounding Box Type	Precision (%)	Recall (%)	mAP@0.5 (%)	F1-Score (%)
Baseline (YOLOv8)	HBB (Horizontal)	82.5	80.2	81.4	81.3
RoI Transformer	OBB (Oriented)	85.3	83.1	84.2	84.2
Oriented R-CNN	OBB (Oriented)	88.4	86.7	87.8	87.5
YOLOv8-OBB (Standard)	OBB (Oriented)	89.2	87.4	88.6	88.3
Proposed Method	OBB (Oriented)	92.8	90.1	91.6	91.4

Note: Bold values indicate the best performance among all compared detection models.

As shown in [Table 5](#), the conventional HBB-based baseline achieved relatively limited performance on the PierCrack-OBB dataset, with an mAP@0.5 of 81.4%. This indicates that when crack targets are slender, curved, and densely distributed, the HBB-based representation is not sufficiently suited to the geometric characteristics of bridge pier cracks. A major reason is that HBBs cannot tightly fit the actual crack contours and often include a substantial proportion of irrelevant background, thereby increasing the model's sensitivity to non-crack textures and surface noise [19]. In contrast, OBB-based methods achieved better performance across all evaluation metrics, suggesting that oriented bounding boxes are more suitable for representing bridge pier cracks, which typically exhibit high aspect ratios. After further integrating DCNv3, a channel-spatial attention mechanism, and Gaussian probability loss into the standard YOLOv8-OBB framework, the proposed method improved the mAP@0.5 to 91.6%, while Precision and Recall reached 92.8% and 90.1%, respectively. Although mAP@0.5 was used as the main indicator for comparison with common object-detection studies, stricter IoU thresholds were also evaluated to assess

localization quality. The proposed method achieved an mAP@0.75 of 84.7% and an mAP@0.5:0.95 of 64.3% on the test set. Compared with the standard YOLOv8-OBB model, the proposed method still maintained higher performance under these stricter criteria, indicating that the improvements enhanced not only target-level detection but also rotated-box localization accuracy. These results indicate that the proposed improvements are beneficial for feature extraction, background suppression, and oriented box regression.

To further analyze the contribution of each improvement module, an ablation study was conducted, and the results are summarized in Table 6.

Table 6: Ablation study of the proposed improved modules on detection performance.

Experimental Group	Baseline (YOLOv8-OBB)	DCNv3	Channel-Spatial Attention	Gaussian Probability Loss	mAP@0.5 (%)	Gain
1	✓				88.6	–
2	✓	✓			90.2	+1.6
3	✓		✓		89.5	+0.9
4	✓			✓	89.8	+1.2
5 (Proposed)	✓	✓	✓	✓	91.6	+3.0

Table 6 shows that each module improves detection performance to a certain extent. Incorporating DCNv3 increases mAP@0.5 by 1.6%, suggesting that deformable convolution enhances the representation of curved crack morphology and locally non-rigid edges. The channel-spatial attention mechanism improves mAP@0.5 by 0.9%, indicating its ability to suppress irrelevant responses in complex backgrounds. Gaussian probability loss increases mAP@0.5 by 1.2%, reflecting its stabilizing effect on angle regression for slender rotated targets. When all three modules are jointly incorporated, the overall performance gain reaches 3.0% over the baseline, demonstrating that these modules have complementary effects in crack feature extraction, background suppression, and oriented box regression. To reduce the influence of training randomness, the random seed was fixed in all experiments, and all ablation variants were trained and evaluated under the same experimental settings. Therefore, the ablation results are interpreted as empirical performance comparisons under a consistent experimental protocol rather than as formal statistical significance tests. The proposed full model achieved the highest mAP@0.5 among all compared variants, indicating that the three introduced modules provide complementary improvements in crack feature extraction, background suppression, and oriented box regression. Nevertheless, further repeated training with multiple random seeds will be considered in future work to provide a more rigorous statistical evaluation.

To further assess model behavior under different confidence thresholds, Precision-Recall (P-R) curves of the conventional HBB model and the improved OBB model were plotted on the test set, as shown in Fig. 6. To improve visual readability, the P-R curves were redrawn using high-resolution output, enlarged axis labels, unified line widths, and clearer legends. The P-R curve comprehensively reflects the trade-off between Precision and Recall under varying detection thresholds [20]. It can be observed from Fig. 6 that the curve of the proposed improved OBB model consistently lies above that of the HBB baseline, with a more gradual decline. This indicates that the proposed model can maintain a relatively high Precision while preserving a high Recall level. The results suggest that, by adopting a more suitable oriented representation for slender cracks together with corresponding network improvements, the detection stability in complex bridge pier scenes is significantly enhanced.




Figure 6: Precision–recall curve comparison between the HBB baseline and the proposed improved OBB model on the bridge pier crack test set.

In addition, considering that bridge pier cracks are often densely distributed in local regions, the false positive and false negative distributions of the two models were further analyzed statistically, as shown in Fig. 7. The error-statistics plot was regenerated with clearer axis labels, consistent font size, direct numerical annotations, and improved top-margin spacing to make the false-positive and false-negative distributions easier to compare. In dense crack regions, the conventional HBB model tends to generate substantial overlap among bounding boxes, making adjacent targets more likely to be incorrectly suppressed during non-maximum suppression, which leads to missed detections. Meanwhile, large background redundancy may also increase the risk of false alarms. In contrast, the improved OBB model can more tightly fit crack contours and therefore achieves better separation of dense targets. The quantitative comparison is provided in Table 7.

As indicated in Table 7, in dense crack scenarios, the proposed method achieves a substantially lower missed detection rate and significantly reduces the proportion of background redundancy compared with the HBB baseline, owing to the use of oriented representation together with improved feature extraction and regression strategies. In addition, for extremely slender cracks with an aspect ratio greater than 20, the proposed method achieves a recall of 89.2%, which is markedly higher than that of the HBB model. These results further indicate that the proposed method is suitable for detecting slender and densely distributed cracks on bridge piers.

3.2 Visual Comparison of Geometric Representation Capability

In addition to quantitative metrics, the quality of geometric representation in detection results is another important aspect for evaluating engineering applicability. For bridge piers, the spatial distribution, local orientation, and density of cracks are important references for subsequent condition assessment and manual inspection [21]. Therefore, a visual comparison between HBB- and OBB-based representations was further conducted.




Figure 7: Comparison of false-positive and false-negative error statistics between the HBB baseline and the proposed improved OBB model in dense crack scenarios.

Table 7: Quantitative analysis of missed detection rates and background redundancy area proportions of different detection methods in dense crack scenarios.

Detection Model	Missed Detection Rate due to NMS in Dense Areas (%)	Avg. Background Redundancy Area Proportion within Bbox (%)	Recall Rate for Extremely Slender Targets (Aspect Ratio > 20) (%)
Baseline (YOLOv8-HBB)	28.40%	91.20%	45.60%
Proposed (Improved OBB)	3.50%	12.40%	89.20%

Note: The recall rate for extremely slender targets refers to the recall of crack samples with an aspect ratio greater than 20, and is used to evaluate the model's capability in detecting very slender cracks.

Fig. 8 presents representative detection results in a real bridge pier scene with densely distributed vertical cracks. The original real detection images were retained, and the revised figure was exported at higher resolution with clearer panel labels, enhanced contrast, and more distinguishable HBB and OBB annotations to better illustrate the difference in geometric representation capability. Fig. 8a shows the output of the conventional HBB model. It can be observed that in regions where cracks are slender and closely adjacent, horizontal bounding boxes are unable to align well with the actual crack contours, and phenomena such as one box covering multiple cracks or severe overlap among neighboring boxes frequently occur. Meanwhile, a large amount of concrete background is included within the boxes, weakening the ability of the detection results to faithfully represent the true geometric boundaries of cracks. Fig. 8b shows the results of the proposed improved OBB model. Compared with HBBs, OBBs can more closely fit the local orientation and contour morphology of cracks, enabling clearer separation of adjacent cracks in dense regions while reducing the inclusion of irrelevant background. These observations indicate that oriented representation is more appropriate for the geometric description of high-aspect-ratio crack targets on bridge piers.




Figure 8: Visual comparison of geometric crack representation in a dense vertical crack scenario: (a) HBB detection with background redundancy and box overlap; (b) multi-segment OBB detection with improved crack-contour fitting and adjacent-crack separation.

To further analyze the differences in feature attention between the two models, Grad-CAM activation heatmaps were generated for the corresponding detection regions, as shown in Fig. 9. The original Grad-CAM results were retained and re-exported at higher resolution with a clearer panel arrangement, enlarged labels, and a unified color scale to better compare the feature activation patterns of the HBB and OBB models. For the HBB model, high-response regions are relatively dispersed, and strong activations are observed not only on the crack areas but also on weathered textures, water stains, and illumination variations on the pier surface, indicating that the model is affected to some extent by background interference [22]. In contrast, the high-response regions of the proposed improved OBB model are more concentrated on the crack body and its continuous propagation path, indicating that the model can focus more effectively on crack-related features. This observation is consistent with the quantitative results presented above and further supports the improved crack feature extraction and geometric representation capability of the proposed method under complex backgrounds.

For difficult inspection conditions, such as low illumination, rain stains, efflorescence, algae coverage, surface roughness, and slight motion blur, the proposed method generally maintained better crack localization than the HBB baseline when the crack boundary remained visually identifiable. This robustness can be attributed to the adaptive sampling ability of DCNv3, the background-suppression effect of the channel-spatial attention module, and the tighter geometric representation of OBBs. However, severe motion blur, heavy algae or stains covering the crack body, and extremely low contrast may still reduce detection confidence or cause inaccurate minor-axis estimation. Therefore, under such severe conditions, the proposed method is more suitable for preliminary screening, and suspicious regions should still be verified by close-range inspection or higher-quality image acquisition.

Overall, the above quantitative results and visual comparisons show that the proposed improved OBB model not only outperforms the conventional HBB method in detection accuracy, but also provides better performance in dense crack separation, background redundancy control, and local crack orientation

representation. These advantages offer a more reliable basis for subsequent crack width quantification and geometric orientation analysis.




Figure 9: Grad-CAM visualization of feature activation under weathered bridge pier backgrounds: (a) dispersed activation of the HBB model; (b) crack-focused activation of the proposed improved OBB model.

In addition to detection accuracy, the computational efficiency of the proposed model was also considered to evaluate its engineering applicability in UAV-assisted bridge inspection. Under the same hardware environment described above and using an input size of 640×640 pixels, the proposed model contains approximately 13.1 M parameters and achieves an average inference time of 11.8 ms per image on the NVIDIA RTX A6000 GPU, corresponding to about 84.7 FPS. Compared with the standard YOLOv8-OBB model, the introduction of DCNv3, the channel-spatial attention module, and Gaussian probability loss slightly increases the computational cost, but the additional processing time remains limited. Considering that the mAP@0.5 is improved from 88.6% to 91.6%, the proposed method achieves a reasonable balance between detection accuracy and inference efficiency. Therefore, it is suitable for offline batch processing of UAV inspection images and has potential for near-real-time auxiliary screening in practical bridge pier crack inspection.

4 Quantification of Crack Geometric Features and Engineering Analysis

4.1 Automated Quantification of Maximum Effective Crack Width

In bridge structural health monitoring, automatic crack detection is only the first step of damage assessment. A key challenge for practical intelligent inspection lies in converting detection results into physically meaningful quantities with engineering relevance. Instead of providing only pixel-level locations, the proposed method further exploits the geometric parameters of oriented bounding boxes (OBBs) to establish a mapping from image-based detection results to physical crack width. Specifically, based on the ground sample distance (GSD) defined in Section 2, the minor-axis pixel length of each detected OBB is converted into the corresponding physical dimension, thereby enabling automatic estimation of the local transverse opening width of the crack. Since the annotation protocol specifies that the major axis of the OBB follows the crack direction while the minor axis aligns with the width-normal direction, the minor axis can be used to approximate the local width characteristic of the corresponding crack segment.

To verify the effectiveness of this quantification method, the estimated crack widths were compared with manual field measurements. The ground-truth crack width was established from physical field measurements rather than from image annotation alone. A total of 626 representative crack-width measurement samples were selected from the test images, covering different width intervals and crack morphologies. Manual measurements were conducted on site using an optical crack microscope and feeler gauges, with the measurement resolution recorded according to the instrument specifications. For each selected crack segment, the measurement point was chosen at the visually identifiable maximum opening position within the corresponding local crack segment, and repeated readings were averaged to reduce random measurement error. The manually measured location was then matched with the detected OBB segment according to the image position and local crack path; the algorithm-estimated width was calculated by multiplying the minor-axis length of the matched OBB by the corresponding GSD value. Therefore, the averaged physical measurement at each matched crack segment was used as the ground-truth reference value for calculating MAE and RMSE. As shown in Fig. 10, a generally consistent relationship can be observed between the algorithm-estimated widths and the manually measured values, indicating that the proposed width quantification strategy based on the OBB minor axis and GSD is feasible. The crack-width regression plot was regenerated with clearer scatter points, an ideal 1:1 reference line, a ± 0.18 mm error-envelope boundary, the 0.20 mm engineering threshold line, and quantitative error indicators. The revised figure reports $N = 626$, $MAE = 0.154$ mm, and $RMSE = 0.195$ mm, which are consistent with the overall error level reported in the manuscript. Statistical results show that the mean absolute error (MAE) of the proposed method is less than 0.18 mm over the full sample range, demonstrating its capability for relatively stable millimeter-level crack width estimation [23].




Figure 10: Regression comparison between algorithm-estimated crack widths and manual field-measured reference values.

From the perspective of engineering applications, crack width is an important indicator in the durability assessment and maintenance decision-making of concrete bridge components. According to the crack-width control requirements in the Specifications for Design of Highway Reinforced Concrete and Prestressed Concrete Bridges and Culverts (JTG 3362—2018), crack-width checking is required for reinforced concrete and partially prestressed concrete bridge members under serviceability limit states [24]. In this study, the 0.2 mm value was used as an engineering screening threshold for crack-width classification and algorithm-performance evaluation, rather than as a universally applicable safety limit. The allowable crack width may vary with structural type, reinforcement condition, exposure environment, durability requirement, and local maintenance specification. Therefore, cracks close to or exceeding this threshold should be further interpreted together with exposure conditions and verified through field inspection or structural assessment.

To further evaluate the quantification performance under different width intervals, the error statistics for each interval were summarized in Table 8. It can be seen that the proposed method maintains good estimation performance across different crack width ranges. In particular, within the critical interval of 0.15–0.25 mm, the model still achieves relatively high classification accuracy. This indicates that the proposed method retains certain sensitivity near engineering threshold values and may support automated screening based on code-specified limits.

Table 8: Error statistics of crack-width estimation under different physical width intervals.

Crack Width Interval (mm)	Number of Samples	Mean Absolute Error, MAE (mm)	Root Mean Square Error, RMSE (mm)	Maximum Signed Estimation Error (mm)	Classification Accuracy for the 0.2 mm Screening Threshold (%)
0.05–0.15 (Micro)	215	0.11	0.14	0.28	98.2
0.15–0.25 (Critical Zone)	184	0.14	0.18	0.31	95.6
0.25–0.50 (Distinct)	142	0.18	0.22	−0.35	100.0
>0.50 (Large-width cracks)	85	0.25	0.31	−0.42	100.0
Overall	626	0.16	0.20	−0.42	97.8

Note: In the maximum estimation error, positive values indicate overestimation and negative values indicate underestimation. The classification accuracy for the 0.2 mm screening threshold refers to the accuracy of determining whether the estimated crack width exceeds the engineering screening threshold adopted in this study. This threshold is used for comparative evaluation and preliminary screening, and should not be interpreted as a universal allowable limit for all bridge types, exposure environments, or durability conditions.

The crack width estimation results are influenced by several sources of uncertainty related to image acquisition, geometric conversion, crack morphology, and reference measurement. First, shooting-distance deviation and camera-pose variation may change the local imaging scale, especially when the optical axis of the UAV camera is not approximately perpendicular to the pier surface. Under oblique, upward-looking, or downward-looking views, perspective distortion may make the GSD calculated from the nominal shooting distance inconsistent with the actual local scale along the crack-width direction, which may lead to overestimation or underestimation of the OBB minor-axis-based width. Although images with severe perspective distortion or uncertain local surface orientation were excluded from the primary physical width validation samples, residual scale errors may still remain when the crack region lies on a curved surface, when the local surface normal is difficult to estimate, or when the image contains strong perspective deformation. Second, bridge pier surface unevenness and curvature may cause local depth variations within the same image, especially for cylindrical or rounded piers, resulting in slight deviations in the pixel-to-millimeter conversion. Third, residual lens distortion after camera calibration may also introduce local geometric deformation near image edges, although this effect is relatively limited when the crack region is close to the image center. In addition, illumination changes, shadows, water stains, efflorescence, surface roughness, and

weathered concrete textures may reduce the contrast between the crack and the surrounding background, blur the visible crack boundary, and affect the predicted minor-axis length of the oriented box. This effect is more pronounced for irregular, branched, rough-edged, or partially shadowed cracks, for which the OBB minor-axis-based width should be regarded as an approximate effective width and should be verified by close-range inspection when the estimated value is close to an engineering threshold. Finally, the manual field measurements used as reference values may also contain errors caused by instrument precision, subjective selection of measurement points, and the spatial mismatch between the physical measurement position and the corresponding detected OBB segment. Therefore, the reported width estimation accuracy should be interpreted as the performance under the adopted UAV imaging and field measurement conditions [23]. In practical bridge inspection, the proposed method is more suitable for rapid screening, width-level classification, and prioritization of suspicious cracks, while cracks close to maintenance thresholds or located in critical structural regions should still be verified by close-range inspection or high-precision contact measurement.

The proposed quantification method is based on 2D image measurements and therefore cannot fully recover the three-dimensional opening morphology of cracks. Surface curvature, local depth irregularity, uneven illumination, camera defocus, motion blur, and oblique viewing angles may change the apparent crack boundary and local imaging scale, leading to uncertainty in the OBB minor-axis-based width estimation. In particular, when the crack edge is partially shadowed, out of focus, or located on a curved pier surface, the estimated width should be regarded as an approximate visible surface width rather than the true internal crack opening. Therefore, the proposed 2D image-based quantification is suitable for rapid screening and comparative width-level assessment, but critical cracks still require close-range measurement or three-dimensional/field verification for reliable engineering diagnosis.

Overall, the proposed OBB- and GSD-based quantification method converts pixel-level detection outputs into physical-scale indicators, providing a basis for engineering-oriented interpretation of crack detection results. This method shows potential for rapid crack width screening and quantitative statistical analysis in bridge pier inspection.

4.2 Analysis of Crack Geometric Morphology and Orientation Characteristics

In addition to crack width, crack morphology and spatial orientation are also important sources of information in bridge defect assessment. For bridge piers, features such as aspect ratio, propagation direction, and overall spatial distribution not only describe the geometric properties of surface damage, but may also provide supportive information for structural condition analysis. Therefore, based on the detection results, the geometric morphology and orientation distribution of cracks were further analyzed in this study.

First, in terms of geometric morphology, the crack samples in the PierCrack-OBB dataset exhibit pronounced high-aspect-ratio characteristics. Fig. 11 presents the histogram of crack aspect ratio distribution for bridge piers. The aspect-ratio histogram was redrawn with unified fonts, clearer axis labels, and statistical annotations to better demonstrate the long-tail distribution of slender crack targets. As shown in Fig. 11, the distribution is strongly right-skewed with a long-tail pattern, and most cracks belong to slender targets. Statistical analysis shows that more than 86.1% of the cracks have aspect ratios greater than 10, with an average aspect ratio of 26.4 and a maximum aspect ratio of 82.5. These results indicate that bridge pier cracks are typically slender and highly elongated targets, further supporting the necessity of using oriented bounding boxes for their representation. Compared with conventional horizontal bounding boxes, OBBs are better suited to accommodate the aspect ratio characteristics and local directional variations of such targets.




Figure 11: Aspect-ratio distribution of bridge pier cracks in the PierCrack-OBB dataset.

Second, based on the inclination angle output by the OBB detector, the spatial orientation of cracks on bridge piers was statistically analyzed, and the resulting rose diagram is shown in Fig. 12. For visualization in the rose diagram, the axial crack orientations within 0° – 180° were mirrored to 180° – 360° , so that opposite directions represent the same crack orientation. The crack orientation rose diagram was regenerated with clearer angular labels, unified radial grids, improved visual contrast, and an explanatory note for axial orientation visualization. As shown in Fig. 12, the orientation distribution is not uniformly dispersed but exhibits clustering within several angular intervals, with a dominant vertical orientation and secondary diagonal orientations. This indicates that the angle information provided by oriented bounding boxes can support crack orientation characterization and geometric statistical analysis.

The orientation-based categories in this study are not intended to serve as independently validated structural diagnostic conclusions. They were established as auxiliary interpretation categories based on commonly observed crack-orientation patterns in concrete bridge components, the visible field morphology of the detected cracks, and the geometric inclination information provided by OBBs. No direct structural loading test or finite-element back-analysis was performed to prove a one-to-one correspondence between each orientation category and a specific mechanical mechanism. Therefore, the terms such as vertical compression-related cracks and diagonal shear-related cracks are used as possible engineering interpretations rather than definitive failure-mode diagnoses.

To facilitate engineering interpretation of the above statistical results, the inclination distribution characteristics were further compared with the crack orientation classification criteria, and the results are summarized in Table 9. As shown in Table 9, the crack inclination angles in the test set exhibit a dominant peak in the interval of 85° – 95° , corresponding to predominantly vertically extending crack patterns. In addition, secondary peaks can be observed around 40° – 50° and 130° – 140° , indicating that some cracks also present diagonal distribution characteristics. Combined with common crack patterns observed in bridge structures, these results indicate that the surface cracks on bridge piers include predominantly vertical cracks and a smaller number of diagonal cracks. These orientation patterns may provide clues for possible

compression-related splitting, local shear effects, or more complex stress conditions; however, the underlying mechanical causes cannot be determined solely from surface crack orientation. Additional evidence from structural loading conditions, reinforcement layout, material deterioration, and field inspection is required before a definite engineering diagnosis can be made [25].




Figure 12: Rose diagram of crack orientation distribution based on OBB inclination angles.

Table 9: Orientation-based interpretation of macroscopic crack patterns and possible associated mechanical implications.

OBB Inclination Peak Interval	Corresponding Classification Threshold (Table 4)	Macroscopic Spatial Orientation	Possible Mechanical Implication	Auxiliary Engineering Interpretation
85°–95° (Absolute Dominant)	80°–100° (Vertical crack pattern)	Vertical Extension	May be associated with compression-related splitting or transverse tensile strain caused by Poisson's effect.	The pier may be affected by axial compression, eccentric loading, or uneven foundation settlement; this interpretation should be verified through structural inspection and field diagnosis.

(Continued)

Table 9 (continued)

OBB Inclination Peak Interval	Corresponding Classification Threshold (Table 4)	Macroscopic Spatial Orientation	Possible Mechanical Implication	Auxiliary Engineering Interpretation
40°–50°/130°–140° (Secondary Peak)	30°–60°/120°–150° (Diagonal crack pattern)	Diagonal Distribution	May be associated with local shear-related stress concentration or complex stress conditions.	Localized areas of the pier may be affected by lateral actions such as wind loads, water impact, or secondary torsion; further structural evidence is required for confirmation.
0°–15°/165°–180° (Sparse Distribution)	0°–15°/165°–180° (Horizontal/local flexural crack pattern)	Horizontal Extension	May be associated with local flexural effects or tensile stress concentration in the edge region.	The pier may be affected by lateral bending effects, and further field verification is required.

The proposed method does not distinguish structural cracks from non-structural surface cracks solely based on geometric orientation. The OBB inclination angle provides useful information for describing crack direction, but structural relevance should be judged together with crack width, aspect ratio, spatial continuity, distribution density, location on the pier, and visible surface conditions. Randomly distributed, map-like, shallow, or texture-associated cracks were treated as possible non-structural surface cracks only when their morphology and field appearance supported such an interpretation. Conversely, vertical or diagonal cracks with larger width, stronger continuity, and concentrated spatial distribution were regarded as cracks requiring further structural review rather than being automatically classified as structural cracks. Therefore, orientation-based categories should be used only as auxiliary descriptors for crack morphology and supporting information for subsequent engineering review. A reliable structural diagnosis still requires integration with field inspection, structural drawings, load analysis, material testing, and, when necessary, numerical or mechanical verification.

In engineering practice, the proposed crack orientation classification is mainly applicable to rapid geometric screening and preliminary defect organization. For example, it can be used to summarize the dominant propagation direction of cracks in a bridge pier, identify regions where vertical or diagonal cracks are concentrated, assist inspectors in prioritizing suspicious areas, and provide structured geometric information for subsequent maintenance records. However, the classification results should be regarded as auxiliary indicators rather than independent diagnostic conclusions. Their applicability is limited by image quality, UAV viewing angle, perspective distortion, pier surface curvature, crack branching, and the fact that different mechanical or environmental causes may produce similar surface crack orientations. Therefore, the orientation classification should be interpreted together with crack width, crack density, spatial location, structural drawings, loading conditions, historical inspection records, and on-site verification before being used in maintenance decision-making.

Overall, on the basis of crack detection, the proposed framework further utilizes the minor-axis and inclination-angle parameters of OBBs to perform crack width estimation and orientation statistical analysis, thus establishing a conversion path from image-based detection outputs to quantifiable geometric features. This extends the role of conventional object detection results in bridge defect assessment: the detection output can not only indicate crack location, but also support crack width statistics, orientation characterization, and auxiliary engineering analysis. Therefore, the proposed method provides a useful reference for the refined identification and quantitative evaluation of bridge pier cracks. Compared with pixel-level semantic

segmentation methods, which can provide detailed crack masks but usually require dense annotations and may be sensitive to ambiguous crack boundaries, the proposed OBB-based strategy emphasizes a practical balance among annotation efficiency, detection robustness, geometric representation, and physical quantification capability [26].

5 Conclusions

To address the limitations of conventional computer vision methods in bridge pier crack characterization, including geometric mismatch, background redundancy, and insufficient expression of key geometric information, this study proposed a high-fidelity characterization and physical quantification method for bridge pier cracks based on oriented object detection (OBB). By constructing a dedicated PierCrack-OBB dataset from real service environments and improving the YOLOv8-OBB framework with DCNv3, a channel-spatial attention mechanism, and Gaussian probability loss, a technical pipeline was established from crack detection and geometric representation to physical quantification. Based on the results of this study, the following conclusions can be drawn.

First, the applicability and advantages of oriented object detection for representing slender cracks on bridge piers were verified. Compared with conventional horizontal bounding box (HBB)-based methods, OBBs can better fit the slender morphology and local directional variation of cracks, reduce background redundancy, and improve target separation in dense crack scenarios. Experimental results show that the proposed improved model achieved an mAP@0.5 of 91.6% on the test set, demonstrating better detection accuracy and stability than the HBB baseline. These findings indicate that representing bridge pier cracks with oriented bounding boxes helps improve the geometric fidelity and detection reliability of crack identification under complex backgrounds.

Second, a quantitative mapping relationship was established from the geometric parameters of oriented bounding boxes to the physical features of cracks. The results show that, by combining the minor-axis parameter of OBBs with the ground sample distance (GSD) of UAV inspection images, automated crack width estimation can be achieved with a mean absolute error of less than 0.18 mm, indicating favorable millimeter-level quantification capability. In addition, statistical analysis of OBB inclination angles enables the description of crack orientation characteristics and provides auxiliary information for crack morphology recognition and subsequent engineering review. These results suggest that the output of oriented object detection can be extended beyond crack localization to support quantitative analysis of bridge defects.

Finally, this study shows that integrating oriented object detection with physical crack quantification provides a feasible approach for automatic identification, geometric characterization, and auxiliary assessment of bridge pier cracks. Compared with conventional vision-based methods that mainly remain at the target localization stage, the proposed method further extends the potential value of detection results in crack width statistics, orientation analysis, and engineering applications. The proposed framework can provide technical support for intelligent inspection and refined defect assessment in bridge structural health monitoring. In future work, it is still necessary to incorporate more bridge types and more complex service conditions in order to further evaluate the cross-scene generalization ability of the model, improve the accuracy of physical-scale quantification, and explore deeper integration with structural mechanics analysis methods.

Acknowledgement: Not applicable.

Funding Statement: The authors received no specific funding for this study.

Author Contributions: The authors confirm contribution to the paper as follows: Conceptualization, Yu Li and Xinxin Cao; Methodology, Yu Li; Software, Zongqi Xiong; Validation, Yu Li, Zongqi Xiong and Xinxin Cao; Formal analysis, Yu Li; Investigation, Zongqi Xiong; Resources, Xinxin Cao; Data curation, Zongqi Xiong; Writing—original draft preparation, Yu Li; Writing—review and editing, Xinxin Cao; Visualization, Zongqi Xiong; Supervision, Xinxin Cao. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The data that support the findings of this study are available from the corresponding author, Xinxin Cao, upon reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Marra A, Trizio I, Fabbrocino G. An automated information modeling workflow for existing bridge inspection management. *Infrastructures*. 2025;10(11):312. doi:10.3390/infrastructures10110312.
2. Zhang C, Zou Y, Wang F, del Rey Castillo E, Dimyadi J, Chen L. Towards fully automated unmanned aerial vehicle-enabled bridge inspection: where are we at? *Constr Build Mater*. 2022;347(6):128543. doi:10.1016/j.conbuildmat.2022.128543.
3. Cha YJ, Ali R, Lewis J, Büyüköztürk O. Deep learning-based structural health monitoring. *Autom Constr*. 2024;161(3):105328. doi:10.1016/j.autcon.2024.105328.
4. Dan D, Dan Q. Automatic recognition of surface cracks in bridges based on 2D-APES and mobile machine vision. *Measurement*. 2021;168(6):108429. doi:10.1016/j.measurement.2020.108429.
5. Hsieh YA, Tsai YJ. Machine learning for crack detection: review and model performance comparison. *J Comput Civ Eng*. 2020;34(5):04020038. doi:10.1061/(asce)cp.1943-5487.0000918 Erratum in: *J Comput Civ Eng*. 2021;35(1):08220001. doi:10.1061/(ASCE)CP.1943-5487.0000939.
6. Pennada S, Perry M, McAlorum J, Dow H, Dobie G. Threshold-based BRISQUE-assisted deep learning for enhancing crack detection in concrete structures. *J Imag*. 2023;9(10):218. doi:10.3390/jimaging9100218.
7. Dung CV, Anh LD. Autonomous concrete crack detection using deep fully convolutional neural network. *Autom Constr*. 2019;99(4):52–8. doi:10.1016/j.autcon.2018.11.028.
8. Billah UH, La HM, Tavakkoli A. Deep learning-based feature silencing for accurate concrete crack detection. *Sensors*. 2020;20(16):4403. doi:10.3390/s20164403.
9. Ni F, Zhang J, Chen Z. Pixel-level crack delineation in images with convolutional feature fusion. *Struct Control Health Monit*. 2019;26(1):e2286. doi:10.1002/stc.2286.
10. Li S, Zhao X. Automatic crack detection and measurement of concrete structure using convolutional encoder-decoder network. *IEEE Access*. 2020;8:134602–18. doi:10.1109/ACCESS.2020.3011106.
11. Park SE, Eem SH, Jeon H. Concrete crack detection and quantification using deep learning and structured light. *Constr Build Mater*. 2020;252(5):119096. doi:10.1016/j.conbuildmat.2020.119096.
12. Jang K, An YK, Kim B, Cho S. Automated crack evaluation of a high-rise bridge pier using a ring-type climbing robot. *Comput Aided Civ Infrastruct Eng*. 2021;36(1):14–29. doi:10.1111/mice.12550.
13. Liu Y, Zhou T, Xu J, Hong Y, Pu Q, Wen X. Rotating target detection method of concrete bridge crack based on YOLO v5. *Appl Sci*. 2023;13(20):11118. doi:10.3390/app132011118.
14. Li Y, Zhao H, Gu H, Wei Y, Xiang Z, Wang Y, et al. Automated defect segmentation and quantification in concrete structures via unmanned aerial vehicle-based lightweight deep learning. *Comput Aided Civ Infrastruct Eng*. 2025;40(26):4465–84. doi:10.1111/mice.70062.
15. Li Y, Zhao H, Wei Y, Bao T, Li T, Wang Q, et al. Vision-guided crack identification and size quantification framework for dam underwater concrete structures. *Struct Health Monit*. 2025;24(4):2125–48. doi:10.1177/14759217241287906.
16. Liu Y, Zhou T, Hong Y, Pu Q, Wen X. Neighborhood shortest distance method for concrete crack width detection in images. *Eng Struct*. 2025;326(6):119519. doi:10.1016/j.engstruct.2024.119519 Corrigendum in: *Eng Struct*. 2025;326:119568. doi:10.1016/j.engstruct.2024.119568.

17. Gonthina M, Chamata R, Duppalapudi J, Lute V. Deep CNN-based concrete cracks identification and quantification using image processing techniques. *Asian J Civ Eng.* 2023;24(3):727–40. doi:10.1007/s42107-022-00526-9.
18. Sohaib M, Arif M, Kim JM. Evaluating YOLO models for efficient crack detection in concrete structures using transfer learning. *Buildings.* 2024;14(12):3928. doi:10.3390/buildings14123928.
19. Li G, Yang Y, Wen Y, Li J. CGSW-YOLO enhanced YOLO architecture for automated crack detection in concrete structures. *Symmetry.* 2025;17(6):890. doi:10.3390/sym17060890.
20. Chen B, Zhang H, Wang G, Huo J, Li Y, Li L. Automatic concrete infrastructure crack semantic segmentation using deep learning. *Autom Constr.* 2023;152(2):104950. doi:10.1016/j.autcon.2023.104950.
21. Li H, Zhang H, Zhu H, Gao K, Liang H, Yang J. Automatic crack detection on concrete and asphalt surfaces using semantic segmentation network with hierarchical Transformer. *Eng Struct.* 2024;307:117903. doi:10.1016/j.engstruct.2024.117903.
22. Arbaoui A, Ouahabi A, Jacques S, Hamiane M. Concrete cracks detection and monitoring using deep learning-based multiresolution analysis. *Electronics.* 2021;10(15):1772. doi:10.3390/electronics10151772.
23. Nyathi MA, Bai J, Wilson ID. Deep learning for concrete crack detection and measurement. *Metrology.* 2024;4(1):66–81. doi:10.3390/metrology4010005.
24. JTG 3362-2018. Specifications for design of highway reinforced concrete and prestressed concrete bridges and culverts. Beijing, China: China Communications Press; 2018.
25. Wen L, Cheng Y, Fang Y, Li X. A comprehensive survey of oriented object detection in remote sensing images. *Expert Syst Appl.* 2023;224(4):119960. doi:10.1016/j.eswa.2023.119960.
26. Lee D, Kim J, Lee D. Robust concrete crack detection using deep learning-based semantic segmentation. *Int J Aeronaut Space Sci.* 2019;20(1):287–99. doi:10.1007/s42405-018-0120-5.