

Article

Open Access

Automated optical inspection of FAST's reflector surface using drones and computer vision

Jianan Li^{1*}, Shenwang Jiang¹, Liqiang Song², Peiran Peng¹, Feng Mu¹, Hui Li², Peng Jiang² and Tingfa Xu^{1,3*}

Abstract

The Five-hundred-meter Aperture Spherical radio Telescope (FAST) is the world's largest single-dish radio telescope. Its large reflecting surface achieves unprecedented sensitivity but is prone to damage, such as dents and holes, caused by naturally-occurring falling objects. Hence, the timely and accurate detection of surface defects is crucial for FAST's stable operation. Conventional manual inspection involves human inspectors climbing up and examining the large surface visually, a time-consuming and potentially unreliable process. To accelerate the inspection process and increase its accuracy, this work makes the first step towards automating the inspection of FAST by integrating deep-learning techniques with drone technology. First, a drone flies over the surface along a predetermined route. Since surface defects significantly vary in scale and show high inter-class similarity, directly applying existing deep detectors to detect defects on the drone imagery is highly prone to missing and misidentifying defects. As a remedy, we introduce cross-fusion, a dedicated plug-in operation for deep detectors that enables the adaptive fusion of multi-level features in a point-wise selective fashion, depending on local defect patterns. Consequently, strong semantics and fine-grained details are dynamically fused at different positions to support the accurate detection of defects of various scales and types. Our AI-powered drone-based automated inspection is time-efficient, reliable, and has good accessibility, which guarantees the long-term and stable operation of FAST.

Keywords: FAST, Drone, Deep learning, Feature fusion

Introduction

Observing electromagnetic waves is a core challenge in astronomy. Stars, galaxies, and other astronomical objects emit light from across the spectrum, including the visible as well as radio waves, infrared radiation, *etc.* Earth-based

radio telescopes study radio waves emitted by extra-terrestrial sources. The most familiar type of radio telescope is the steerable paraboloid, a device with a parabolically shaped reflector (the dish), which focuses incoming radio waves onto a small pickup antenna. The radio signals are then amplified to a measurable level. Because naturally occurring radio waves are extremely weak by the time they reach the earth from space, even a cell phone signal is billions of times more powerful than cosmic waves received by telescopes. To improve radio-

Correspondence: Jianan Li (lijianan@bit.edu.cn) or Tingfa Xu (ciom_xtf1@bit.edu.cn)

¹Beijing Institute of Technology, Beijing 100081, China

²National Astronomical Observatories of China, Beijing 100107, China

Full list of author information is available at the end of the article.

These authors contributed equally: Jianan Li, Shenwang Jiang

© The Author(s) 2023



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

wave detection, and thus probe further into space, larger reflector dishes must be built to capture more radio waves.

The Five-hundred-meter Aperture Spherical radio Telescope (FAST), also known as the “China Sky Eye”, is the world’s largest single-dish radio telescope. Its optical geometry is outlined in Fig. 1a. The reflector is a partial sphere of radius $R = 300$ m. The planar partial spherical cap of the reflector has a diameter of 519.6 m, 1.7 times larger than that of the previously largest radio telescope. The large reflecting surface makes FAST the world’s most sensitive radio telescope. It was used by astronomers to observe, for the first time, fast radio bursts in the Milky Way and to identify more than 500 new pulsars, four times the total number of pulsars identified by other telescopes worldwide. More interesting and exotic objects may yet be discovered using FAST.

However, each coin has two sides. A larger reflecting surface is more prone to external damage due to environmental factors. The FAST reflector comprises a total 4,450 spliced trilateral panels (Fig. 1b), made of aluminium with uniform perforations to reduce weight and wind impact. Falling objects (e.g., during the extreme events such as rockfalls, severe windstorms, and hailstorms) may cause severe dents and holes in the panels. Such defects adversely impact the study of small-wavelength radio waves, which demands a perfect dish surface. Any irregularity in the parabola scatters these small waves away from the focus, causing information loss.

The rapid detection of surface defects for timely repair is hence critical for maintaining the normal operation of

FAST. This is traditionally done by direct visual inspection. Skilled inspectors climb up the reflector and visually examine the entire surface, searching for and replacing any panels showing dents and holes. However, this procedure has several limitations. Firstly, there is danger involved in accessing hard-to-reach places high above ground. Secondly, it is labour- and time-consuming to scrutinise all the thousands of panels. Thirdly, the procedure relies heavily on the inspectors’ expertise and is prone to human-based errors and inconsistencies.

The remedy to the shortcomings of manual inspection at FAST is automated inspection. As a first step, we integrated deep-learning techniques with the use of drones to automatically detect defects on the reflector surface. Specifically, we began by manually controlling a drone equipped with a high-resolution RGB camera to fly over the surface along a predetermined route. During the flight, the camera captured and recorded videos of the surface condition. One benefit of the advanced flight stability of drones is that the recorded videos can capture much information on surface details. Moreover, thanks to the GPS device and the RTK module onboard the drone platform, every video frame can be tagged with the corresponding drone location with centimetre-level accuracy. The physical locations of the panels that appear in each frame can thus be determined.

Next, we harnessed deep-learning techniques to detect defects using the captured videos^{1,2}. Computer-vision technology based on deep learning has been widely used for defect detection in concrete and steel structures, such as pavement^{3–5} and bridges^{6,7}. Other studies^{8,9} have detected insulator faults using aerial images of high-voltage transmission lines based on improved YOLOv3 models¹⁰. The automatic optical inspection of large-scale equipment surfaces, such as solar plantations, using aerial imagery is demonstrably highly effective. Vlaminck et al.¹¹ proposed a two-stage approach to the automatic detection of anomalies in large photovoltaic sites using drone-based imaging. Similarly, Tommaso et al.¹² presented a UAV-based inspection system for improved photovoltaic diagnostics based on a multistage architecture built on top of YOLOv3¹⁰. Previous works involving defect detection using aerial imagery were primarily designed to detect large defects and were not reliable for detecting very small defects. In contrast, the present work aims to inspect the large surface of FAST from on high. The defect size is particularly small relative to the surface size. Our aim is therefore to accurately identify and locate surface defects, especially small ones. This introduces new challenges.

An in-depth analysis of the extracted frame imagery revealed two inherent characteristics of FAST’s surface

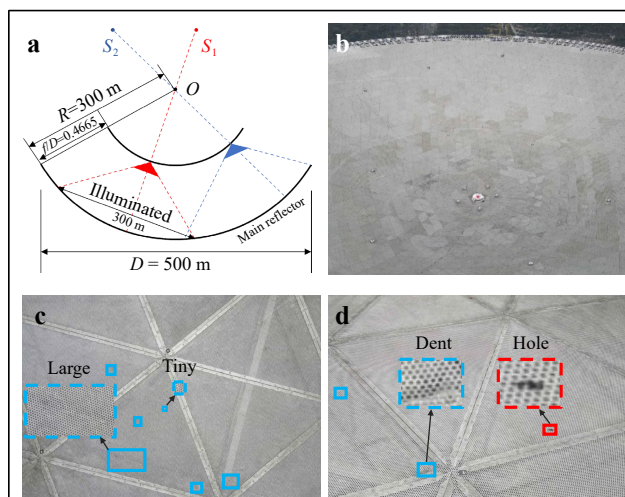


Fig. 1 **a** FAST’s optical geometry. **b** General view of the reflector surface. **c** Large-scale variation of surface defects (dent). **d** High inter-class similarity between different types of defects (dent and hole).

defects: (i) *large-scale variations*, where defects within the same category differ significantly in size (for example, dent sizes can range from 0.4 to 12 inches (Fig. 1c)); (ii) *high inter-class similarity* in visual appearance between defects belonging to different categories (i.e., dents and holes), which makes them difficult to distinguish (Fig. 1d). These unique characteristics strongly increase the challenge of detecting FAST’s surface defects using drone imagery, compared to well-explored general object detection in natural scenes. This therefore requires the introduction of dedicated designs to adapt existing deep detectors to the task and data of interest.

Existing deep-learning methods usually begin by abstracting features from an input image using deep neural networks (DNNs)^{13,14}. This returns hierarchical features with increasing semantics and receptive fields but diminishing details. Most previous general object detectors fuse multi-level features to take their complementary strengths *layer-wise* through addition or concatenation^{15–17}, largely ignoring the spatial distinctiveness within an image. However, in our case, even adjacent positions on the captured drone imagery may have defects of significantly different scales and types. Hence, they require features from different levels to support accurate localisation and recognition. For example, low-level details are preferable in the case of small defects to facilitate localisation, whereas high-level features are better for large or indistinguishable defects to embrace receptive fields and semantics. Hence, multi-level features must be fused in a point-wise selective manner based on defect patterns that vary significantly across the image. This serves to avoid missing detections and generating false positives.

Based on the above considerations, we propose a novel plug-in cross-fusion operation for deep detectors that enables the *point-wise* adaptive fusion of multi-level features. At each position, the features from multiple levels are fused in adaptive proportions, depending on their corresponding content, through a novel point-wise cross-layer attention mechanism. Specifically, cross-fusion defines the feature at each point on a target layer as a *query* and the features at the same point of other layers as *keys*. The *query* feature is refined as a weighted sum of *key* features by modelling their long-range interdependencies. To alleviate the receptive field gap among features from different layers (levels), we further introduced efficient dilated max pooling to enrich low-level features with valuable contextual information. Consequently, the content at each position independently determines which levels of features to aggregate, such that the resulting fused feature map supports the detection of rapidly changing defects across the image. We implemented the cross-fusion

operation as a compact residual learning block that can be flexibly plugged into arbitrary backbones. This improves feature abstraction and, in turn, defect detection with negligible additional parameters and computational cost.

Empowered by the advanced drone platform and dedicated algorithm design, our proposed automated inspection offers clear advantages over conventional manual inspection: (i) *It is safer*, as inspectors need only control the drone remotely from a nearby safe place; (ii) *Accessibility is improved*, as drones can easily reach the upper part of the reflector that is difficult or impossible for humans to reach; (iii) *It is time-efficient*, as the drone’s high mobility allows it to rapidly cover the entire reflector surface, thereby significantly accelerating the inspection process; (iv) *It is highly reliable*, as computer vision is capable of checking hundreds of panels per minute repeatedly, with substantially greater reliability and less error than human inspection.

Methods

Cross Fusion

Given that the defects to be detected vary on a large scale and show high inter-class similarity, we present a novel cross-layer attention mechanism that combines the complementary strengths of multi-level features point by point, depending on the content at each position. Based on the above considerations, an efficient plug-in cross-fusion block was designed to lift the feature representation from the backbone of the detector to support accurate and robust defect detection.

Formulation and Instantiations

Formally, let $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ denote a set of feature maps from multiple levels of the feature hierarchy. Let p represent the index of the target feature map to be refined. The cross-fusion refines the response at point k in \mathbf{x}_p , that is, the *query* feature \mathbf{x}_p^k , as a weighted sum of the responses at point k in multi-level feature maps, namely, *key* feature $\{\hat{\mathbf{x}}_j^k\}_{j=1}^N$.

$$\mathbf{y}_p^k = \sum_{j \in \{1, \dots, N\}} \mathbf{f}(\mathbf{x}_p^k, \hat{\mathbf{x}}_j^k) \phi_j(\hat{\mathbf{x}}_j^k), \quad (1)$$

where $\hat{\mathbf{x}}$ represents a set of reference feature maps obtained by resizing each feature map in \mathbf{x} to the same resolution as \mathbf{x}_p followed by dilated max pooling with a stride of 1. Dilated max pooling is parameter-free and carries negligible additional computational overhead. It is introduced to enlarge the receptive fields of the reference feature maps while retaining their spatial resolution. The settings of the pooling kernel size and dilation rate are discussed in “Extensions to General Object Detection”

Stages 2 to 4, denoted $\{r_2^l, r_3^l, \text{ and } r_4^l\}$, respectively. Considering that Stage 4 contains most residual blocks, we also collect the feature map from its first residual block, denoted r_4^f . We use r_4^l as the target feature map to be refined. We obtain the reference feature maps \hat{x} by downsampling each feature map in x to obtain the same resolution as r_4^l . The block outputs a refined feature map for r_4^l , upon which subsequent feed-forward computations in the backbone are conducted normally.

Data Acquisition

Drone Platform

We used the DJI Mavic 2 Enterprise Advanced (M2EA) drone as the inspection platform, as illustrated in Fig. 3. Equipped with a 48 MP high-resolution RGB camera that supports ultra-zoom, the M2EA can capture small surface details to support accurate defect detection. It features centimetre-level positioning accuracy with the RTK module, which helps to precisely determine the real-world location of any detected surface defect. The light weight of M2EA avoids any damage to the reflector in the event of a flight accident.

Flight Route

The inspectors drove the drone according to a predetermined route, as depicted in the right subfigure of Fig. 3. The route comprised multiple circular paths around

the reflector centre with a gradually increasing radius. The drone flew along every circular path in sequence at a constant speed of 4.0 miles per hour and at a constant vertical height of 1.0 m above the reflector surface. During the flight, distributed measuring bases were used as operating platforms for temporarily recovering and launching the drone (left subfigure of Fig. 4). The well-planned drone flight route allowed a capture of video data covering the entire reflector surface.

Acquiring Data

We divided the surface defects into two classes: *dent* and *hole*. All dents deeper than 0.12 inches and longer than 0.4 inches, and holes with a diameter larger than 0.2 inches were considered as defects to be detected. We extracted 1 out of every 30 frames of the captured videos and reserved the extracted frames containing at least one defect, resulting in a total of 2,582 frames. We used the annotation tool Labellingm to manually label the bounding box and class label for each defect. A total of 34,239 bounding boxes of approximately 3,000 unique defects were labelled. Fig. 5 illustrates the defect size distribution on the collected frames. Both the surface dents and holes show large size variations, which statistically supports our design motivation. We split all labelled data into training and test sets of 2,442 and 140 frames, respectively. The frames in




Fig. 3 Visualisation of the DJI M2EA drone platform (left) and its on-site inspection process (right).

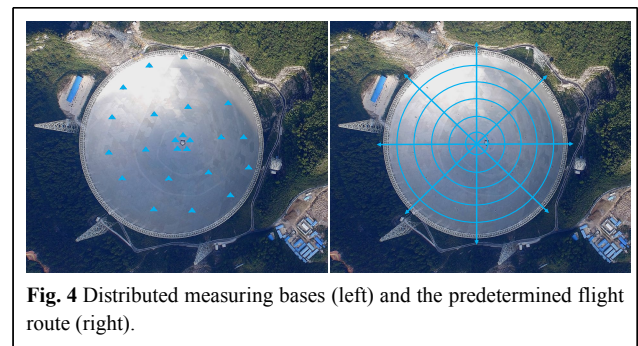


Fig. 4 Distributed measuring bases (left) and the predetermined flight route (right).

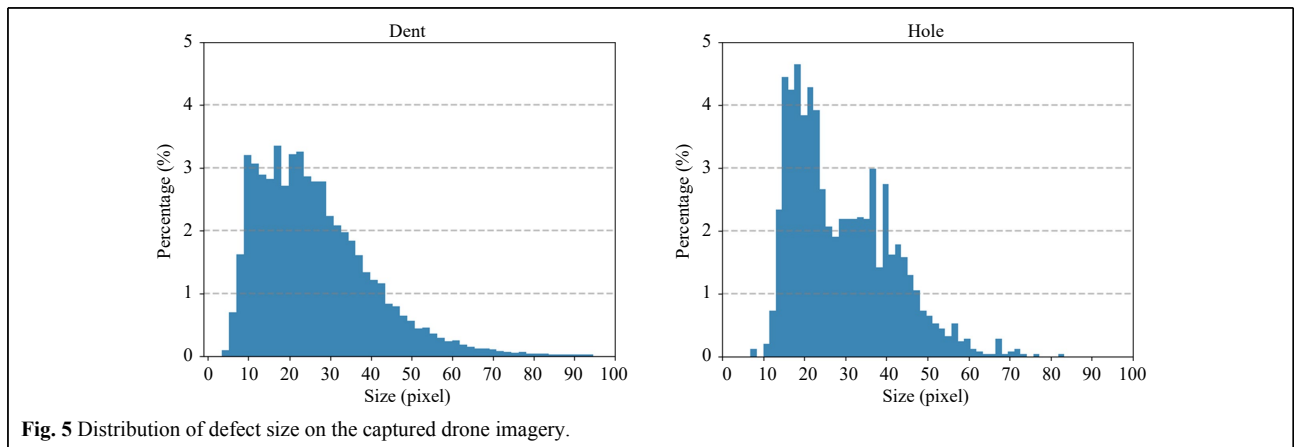


Fig. 5 Distribution of defect size on the captured drone imagery.

these sets cover different areas of the reflector surface, avoiding intersection.

Results

We first report the results of our method for FAST surface defect and pavement distress detection, and then provide evaluations of VisDrone²⁰ and COCO²¹ for drone and general object detection, respectively. Extensive ablation studies were conducted to validate our design choices and parameter settings.

Results for FAST's Surface Defect Detection

Setups. We built our defect detector on top of the recently developed Libra R-CNN²² by incorporating cross-fusion into the ResNet-50 (R50) backbone¹⁴. The weight layers of the inserted cross-fusion block were initialised from a zero-mean Gaussian with a 0.01 standard deviation. We randomly cropped patches of size 1024×1024 from the original frame image and resized them to a maximum of 1333×800 at constant aspect ratio. We trained the detector over 12 epochs using the SGD optimiser with initial learning rate 0.01 and batch size 8. For inference, we extracted patches in a sliding-window manner from the entire frame image. We used COCO-style box Average Precision (AP) and AP with IoU thresholds of 0.50 (AP_{50}) and 0.75 (AP_{75}) as evaluation metrics.

Main results. Table 1 shows that incorporating cross-fusion significantly improves the Libra R-CNN baseline by 1.7% mAP, with 1.2% and 2.1% AP improvements on *dent* and *hole*, respectively. Notably, the stricter criterion of AP_{50} is significantly boosted by 2.4%. This proves that our cross-fusion improves the localisation accuracy, which contributes to the detection of small defects that are critical

to reflector surface inspection. Compared to the YOLO series^{10,23–25} that have been widely used in defect detection on drone imagery^{5,12}, and to the recent efficient TOOD²⁶, our model clearly performs better, in particular surpassing the popular YOLOv5²⁴ by a large margin of 3.4% in the AP. These results confirm the effectiveness of our cross-fusion design.

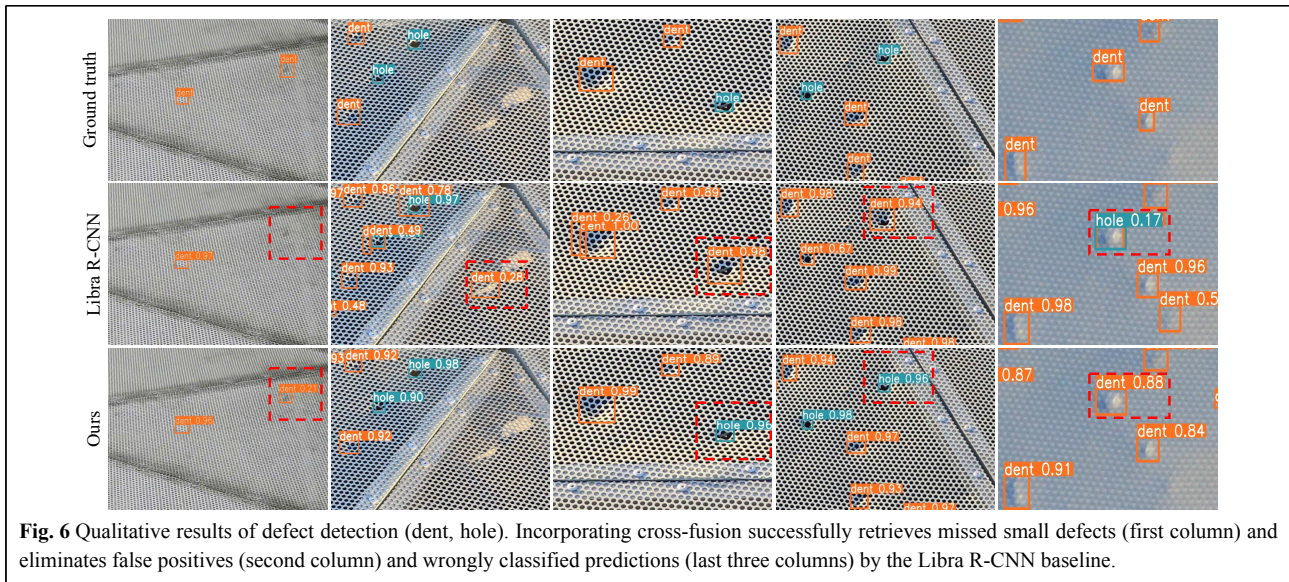
Qualitative results are also shown in Fig. 6. Directly applying Libra R-CNN to defect detection tends to miss small defects (first column) and generate false positives (second column). Moreover, some defects were incorrectly classified because of their similar visual appearance (last three columns). In contrast, incorporating cross-fusion improves the accuracy of bounding boxes for defects at various scales, particularly small ones. Notably, our model successfully detected surface holes as small as 0.2 inches in diameter, which can be easily missed by manual visual inspection. In addition, different types of defects are better distinguished, a critical benefit for instigating timely repairs. These results indicate that our cross-fusion adaptively integrates fine-grained details and strong semantics point-wise depending on the corresponding content, thus enabling accurate defect detection on the entire drone imagery. The results also clearly demonstrate the superiority of our automatic optical inspection over manual inspection, which can be attributed to drone-powered data acquisition and the dedicated cross-fusion design.

Experiments on Pavement Distress Detection

Setups. We tested our method with an automatic optical inspection of pavement. The aim was to detect different types of distress on large-scale pavement with images

Table 1 Results for FAST surface defect detection. AP/ AP_{50} / AP_{75} (%) are reported.

Method	Dent			Hole			AP	AP_{50}	AP_{75}
	AP	AP_{50}	AP_{75}	AP	AP_{50}	AP_{75}			
YOLOv3 ¹⁰	19.8	58.3	8.0	32.4	81.3	19.7	26.1	69.8	13.9
YOLOX ²³	25.4	66.9	13.2	31.2	78.2	18.1	28.3	72.5	15.7
YOLOv5 ²⁴	29.1	74.0	15.2	46.6	91.6	41.8	37.9	82.8	28.5
YOLOv7 ²⁵	31.9	79.8	16.3	45.8	91.4	38.1	38.9	85.6	27.2
TOOD ²⁶	28.8	73.6	14.5	44.1	88.4	34.1	36.4	81.0	24.3
Faster R-CNN ²⁷	29.6	73.4	17.6	45.0	93.3	38.8	37.8	83.4	28.2
CARAFE ²⁸	31.1	76.0	19.6	47.1	93.3	42.8	39.1	84.6	31.2
Grid R-CNN ²⁹	29.5	74.1	15.9	45.9	94.7	41.2	37.7	84.4	28.5
Libra R-CNN ²²	32.3	77.3	19.4	47.0	92.9	40.4	39.6	85.1	29.9
Cross-Pooling (Ours)	33.5	80.0	20.3	49.1	94.9	45.5	41.3	87.5	32.9



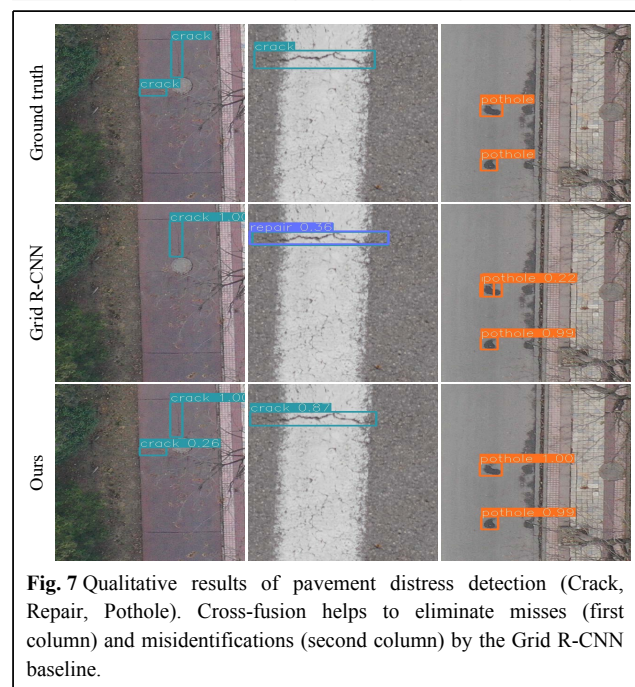
captured from UAVs. We used the pavement image dataset established by Zhu et al.⁵ and divided its publicly available 2,401 images into training and test sets, with proportions of 0.8 and 0.2, respectively. All distress is divided into three categories: *Crack*, *Repair*, and *Pothole*. We built our model on top of the popular Grid R-CNN and trained it over 12 epochs using the SGD optimiser with initial learning rate 0.01 and batch size 8. We report the mean Average Precision (mAP) at (0.5:0.95) IOU.

Main results. As listed in Table 2, our model achieves a mAP of 54.6%, significantly outperforming the Grid R-CNN baseline by 2.3% as well as various well-established detectors, including YOLOv5, TOOD, and CARAFE. In particular, our model improved the Grid R-CNN by 1.3% and 3.5% AP on *Crack* and *Repair*, respectively, indicating its superiority in identifying instances of distress with similar appearance. For the *Pothole* category, where small distress commonly appears, the AP is significantly improved by 1.9%, demonstrating its superior capability of locating small distress caused by cross-fusion.

Qualitative results are also presented in Fig. 7. Our model identifies and locates different types of pavement distress under various scenarios with a high accuracy, demonstrating its clear superiority over previous models. The results reveal that our cross-fusion method benefits defect detection in drone imagery. In addition, applying drones backed up with cross-fusion-based deep models is valuable for the automatic optical inspection of the surfaces of both large-scale equipment and facilities, by providing high accuracy and accessibility while lowering labour costs and time.

Table 2 Results for pavement distress detection. AP_{50:5:95} (%) are reported.

Method	Crack	Repair	Pothole	mAP
YOLOX ²³	27.6	55.5	40.9	41.3
TOOD ²⁶	32.6	55.8	53.2	47.2
Libra R-CNN ²²	36.1	63.9	50.7	50.2
YOLOv5 ²⁴	37.0	64.0	53.2	51.4
CARAFE ²⁸	36.7	60.3	58.5	51.8
Grid R-CNN ²⁹	37.6	62.4	57.0	52.3
Cross-Pooling (Ours)	38.9	65.9	58.9	54.6



Experiments on Drone Object Detection

Setups. Because of the large variance of object scales inherent to drone-derived datasets, we further validated our cross-fusion method using a public drone dataset *VisDrone*²⁰, comprising 10,209 images of 2000×1500 resolution and 10 object categories captured by various drone platforms. We built our detector on top of the recently successful UFPMP-Det³⁰ by simply inserting a single cross-fusion block into the ResNet-50 backbone.

Main results. Table 3 shows our method outperforming the previously best-performing UFPMP-Det. This indicates that incorporating cross-fusion effectively enriches the learned feature representation with strong semantics and fine-grained details, which are crucial for detecting objects at various scales. Fig. 8 provides qualitative results. The UFPMP-Det baseline misses some small-scale objects (the person in the first column) and generates false alarms (last two columns). In contrast, incorporating cross-fusion improves both these results.

Table 3 Results for drone object detection on *VisDrone val*. AP/AP₅₀/AP₇₅ (%) are reported. *: Reproduced results.

Method	AP	AP ₅₀	AP ₇₅
Faster R-CNN ²⁷	21.4	40.7	19.9
ClusDet ³¹	26.7	50.6	24.4
DMNet ³²	28.2	47.6	28.9
GLSAN ³³	30.7	55.4	30.0
AMRNet ³⁴	31.7	52.7	33.1
UFPMP-Det ³⁰	35.8	56.7	38.3
Cross-Pooling (Ours)	36.2	56.9	38.6



Fig. 8 Qualitative results on *VisDrone* (Person, Motor, Car, Pedestrian). Incorporating cross-fusion successfully retrieves missed small objects (indicated by arrows) and eliminates false alarms (indicated by dotted boxes) by the UFPMP-Det baseline.

Extensions to General Object Detection

Setups. Our cross-fusion method can also be generalised for general object detection. We plugged it into well-established general object detectors with various backbones, ranging from ResNet-50/101 (R50/101)¹⁴ to lightweight MobileNetV2 (MNet)¹³. We trained on COCO train2017 and reported the final and ablation results for test2017 and val2017, respectively. The images were resized to a maximum of 1333×800 pixels at constant aspect ratio. The model was trained over 12 epochs using the SGD optimiser with an initial learning rate 0.01 and a batch size 8.

Main results. Table 4 compares the best performing methods. Without resorting to data augmentation and multi-scale testing, incorporating cross-fusion significantly boosted the Faster R-CNN baseline on different backbones by up to 1.7% AP. Our method is similar to recent preminent detectors, such as GRoIE³⁵ and YOLOF³⁶. It

Table 4 Results for general object detection on COCO test2017. AP/AP₅₀/AP₇₅ (%) are reported. *: Higher results reproduced using MMDetection³⁷.

	Model	AP	AP ₅₀	AP ₇₅
MNet	SSDLite ¹³	22.1	–	–
	Faster R-CNN*	23.3	40.7	23.9
	Cross-fusion (Faster R-CNN)	24.4 ↑1.1	42.3	25.0
R50	Faster R-CNN*	35.3	56.3	38.1
	RetinaNet-FPN*	35.9	55.8	38.4
	Faster R-CNN-FPN*	36.6	58.8	39.6
	YOLOF ³⁶	37.5	57.0	40.4
	GRoIE ³⁵	37.5	59.2	40.6
	CARAFE ³⁸	38.1	60.7	41.0
	Libra R-CNN ²²	38.3	59.5	41.9
	Cross-fusion (Faster R-CNN)	37.0 ↑1.7	58.3	39.9
	Cross-fusion (YOLOF)	38.0 ↑0.5	56.9	41.1
	Cross-fusion (Libra R-CNN)	38.6 ↑0.3	59.5	42.2
R101	TDM ⁴¹	35.2	55.3	38.1
	RefineDet ⁴²	36.4	57.5	39.5
	Faster R-CNN*	38.7	59.5	41.8
	M2Det ⁴³	38.8	59.4	41.7
	RetinaNet-FPN ³⁹	39.1	59.1	42.3
	Faster R-CNN-FPN*	39.2	61.2	42.6
	Regionlets ⁴⁰	39.3	59.8	–
	Cross-fusion (Faster R-CNN)	39.5 ↑0.8	60.8	42.7

even surpasses FPN¹⁵, which constructs a feature pyramid with independent predictions made at each level. It also outperforms other fusion-based detectors including TDM⁴¹, RefineDet⁴², and M2Det⁴³. These results clearly demonstrate the superiority of our cross-fusion method in fusing multi-level features for accurate general object detection.

In addition to Faster R-CNN, our method also improves other preeminent detectors, including Libra R-CNN and YOLOF. This evidence indicates that our cross-fusion applies to various detection frameworks, ranging from two-stage Faster R-CNN to one-stage FPN-based Libra R-CNN and one-stage FPN-free YOLOF.

Additional overheads. Our cross-fusion method has a high computational efficiency. We report the additional parameters, FLOPs, and memory footprint introduced by cross-fusion relative to the Faster R-CNN baseline with MobileNetV2 (5.05 M, 96.76 GMac, 1.81 GB), ResNet-50 (33.58 M, 837.07 GMac, 3.92 GB), and ResNet-101 (52.52 M, 913.14 GMac, 4.03 GB). Table 5e shows that a single cross-fusion block introduces negligible additional overhead (e.g., 0.1% FLOPs, 0.4% parameters, and 0.7% memory footprint on MNet) and hardly affects the inference speed.

Impact of correlation function. We tested three instantiations of the pairwise correlation function listed in Table 5a. Generally, cross-fusion with either type of function remarkably improves performance by up to 1.8%

AP. Specifically, the Embedded Gaussian and Sigmoid functions perform similarly and better than the dot product. This indicates that the point-wise dependencies among multi-level features are modelled better by the attention (softmax) or gating (sigmoid) mechanisms. Our experiments used an Embedded Gaussian.

Superiority over conventional fusion strategies. We compared our cross-fusion with conventional multi-level feature fusion strategies, e.g., layer-wise addition, and concatenation. For a fair comparison, we also applied dilated max pooling to these counterpart models. Table 5b shows that our cross-fusion method outperforms both alternative strategies. This indicates that point-wise attentional feature fusion highlights informative features while suppressing less useful ones based on the content at each position, leading to a more powerful feature representation for the entire input image.

Impact of reference features. We investigated the impact of individual reference features by incorporating them incrementally. Table 5c shows that incorporating more reference features (starting from r_3^L) improves performance. Because no improvement is observed by introducing lower-level features, such as r_2^L , we used r_3^L , r_4^F and r_4^L as reference features in our experiments.

Effect of enlarging the receptive field. We introduced dilated max pooling with stride 1 to enlarge the receptive fields of the features. Table 5d shows that removing this operation produces a clear drop in AP. Notably, dilated

Table 5 Ablation studies. Models are trained using Faster R-CNN on COCO train2017, tested on val2017.

(a) Ablation on correlation functions.				(b) Ablation on fusion strategies.				(c) Ablation on reference features.							
model, R50	AP	AP ₅₀	AP ₇₅	model, R50	AP	AP ₅₀	AP ₇₅	model, R50	r_2^L	r_3^L	r_4^F	r_4^L	AP	AP ₅₀	AP ₇₅
baseline	34.9	55.7	37.7	concatenation	36.1	57.0	38.9	res ₄	√				34.9	55.5	37.4
dot product	36.1	57.2	38.8	addition	36.2	57.2	39.1		√				35.1	55.9	38.1
sigmoid	36.6	57.8	39.4	cross-fusion (ours)	36.7	57.7	39.8		√	√			35.7	56.8	38.3
Gaussian	36.7	57.7	39.8						√	√	√		36.7	57.7	39.8

(d) Ablation on the kernel sizes and dilation rates of dilated max pooling.

model, R50	kernel	dilation	AP	AP ₅₀	AP ₇₅
none	–	–	34.7	55.4	37.3
	3×3	3	35.8	56.7	38.7
	5×5	3	36.7	57.7	39.8
max pooling	5×5	2	36.2	57.2	39.1
	5×5	4	36.6	57.9	39.7
convolution	5×5	3	36.7	57.9	39.3

(e) Ablation on backbones. Relative number of parameters and FLOPs are reported. We test inference latency and memory footprint with batch size 1 on a GTX 1080 Ti GPU.

	model	#Params	FLOPs	mem	lat (ms)	AP	AP ₅₀	AP ₇₅
MNet	baseline	1×	1×	1×	114.9	22.9	39.8	23.5
	+ block	1.004×	1.001×	1.007×	116.3	24.1	41.5	24.6
R50	baseline	1×	1×	1×	350.9	34.9	55.7	37.7
	+ block	1.070×	1.011×	1.015×	355.9	36.7	57.7	39.8
R101	baseline	1×	1×	1×	381.7	38.6	59.3	41.6
	+ block	1.050×	1.012×	1.005×	383.1	39.3	60.4	42.4

max pooling performs on par with dilated convolution, while being parameter-free and more computationally efficient. Large kernel sizes and dilation rates are only partly beneficial. We used kernel size 5×5 and dilation rate 3.

Results on different backbones. We tested the backbones at various depths: e.g., MNet and R101. For MNet, we used the bottom 14 convolutional layers as the feature extractor and the remaining layers as the classifier head. We inserted cross-fusion before layer 13 and used the feature maps of layer 7-10 for reference. Table 5e shows cross-fusion improving MNet and R101 by 1.2 and 0.7% of AP, respectively, suggesting that capturing point-wise multi-level dependencies complements the increasing network depth.

Conclusion and Discussions

This work presents an automated optical inspection of the reflector surface of FAST, the world's largest single-dish radio telescope, by exploiting advances in drone technology and deep-learning techniques. To tackle the challenges of surface defects in drone imagery exhibiting large-scale variation and high inter-class similarity, we introduced a simple yet effective cross-fusion operation that aggregates multi-level features in a point-wise selective manner to help detect defects of various scales and types.

Our cross-fusion method is lightweight and computationally efficient, particularly valuable features for onboard drone applications. Currently, we process the video data captured by the drone camera offline on a ground station, which increases the operational complexity. Future work will implement the algorithm on embedded hardware platforms to process captured videos onboard the drone, to make the inspection system more autonomous and more robust.

Acknowledgements

This work was financially supported by the National Natural Science Foundation of China (No. 62101032), the Postdoctoral Science Foundation of China (Nos. 2021M690015, 2022T150050), and the Beijing Institute of Technology Research Fund Program for Young Scholars (No. 3040011182111).

Author details

¹Beijing Institute of Technology, Beijing 100081, China. ²National Astronomical Observatories of China, Beijing 100107, China. ³Beijing Institute of Technology Chongqing Innovation Center, Chongqing, 401135, China

Author contributions

Li, Xu and Song proposed the idea and initiated the project. Li and Jiang performed the theoretical analysis and experimental tests. Peng and Mu collected and analysed the experimental data. J. Li and S. Jiang

wrote the manuscript. T. Xu and L. Song supervised the study.

Conflict of interest

The authors declare that they have no conflicts of interest.

Received: 29 August 2021 Revised: 07 December 2022 Accepted: 08 December 2022

Accepted article preview online: 09 December 2022

Published online: 05 January 2023

References

- Kuschmierz, R. et al. Ultra-thin 3D lensless fiber endoscopy using diffractive optical elements and deep neural networks. *Light: Advanced Manufacturing* **2**, 30 (2021).
- Situ, G. H. Deep holography. *Light: Advanced Manufacturing* **3**, 8 (2022).
- Cao, W. M., Liu, Q. F. & He, Z. Q. Review of pavement defect detection methods. *IEEE Access* **8**, 14531-14544 (2020).
- Cao, M. T. et al. Survey on performance of deep learning models for detecting road damages using multiple dashcam image resources. *Advanced Engineering Informatics* **46**, 101182 (2020).
- Zhu, J. Q. et al. Pavement distress detection using convolutional neural networks with images captured via UAV. *Automation in Construction* **133**, 103991 (2022).
- Zhang, C. B., Chang, C. C. & Jamshidi, M. Concrete bridge surface damage detection using a single-stage detector. *Computer-Aided Civil and Infrastructure Engineering* **35**, 389-409 (2020).
- Du, F. J., Jiao, S. J. & Chu, K. L. Application research of bridge damage detection based on the improved lightweight convolutional neural network model. *Applied Sciences* **12**, 6225 (2022).
- Liu, C. Y. et al. Insulator faults detection in aerial images from high-voltage transmission lines based on deep learning model. *Applied Sciences* **11**, 4647 (2021).
- Liu, J. J. et al. An improved method based on deep learning for insulator fault detection in diverse aerial images. *Energies* **14**, 4365 (2021).
- Redmon, J. & Farhadi, A. YOLOV3: an incremental improvement. Preprint at <https://doi.org/10.48550/arXiv.1804.02767> (2018).
- Vlaminck, M. et al. Region-based CNN for anomaly detection in PV power plants using aerial imagery. *Sensors* **22**, 1244 (2022).
- Di Tommaso, A. et al. A multi-stage model based on YOLOv3 for defect detection in PV panels based on IR and visible imaging by unmanned aerial vehicle. *Renewable Energy* **193**, 941-962 (2022).
- Sandler, M. et al. MobileNetV2: inverted residuals and linear bottlenecks. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Salt Lake City, UT, USA: IEEE, 2018, 4510-4520.
- He, K. M. et al. Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. Las Vegas, NV, USA: IEEE, 2016, 770-778.
- Lin, T. Y. et al. Feature pyramid networks for object detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. Honolulu, HI, USA: IEEE, 2017, 2117-2125.
- Bell, S. et al. Inside-outside net: detecting objects in context with skip pooling and recurrent neural networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. Las Vegas, NV, USA: IEEE, 2016, 2874-2883.
- Chen, L. C. et al. Encoder-decoder with atrous separable convolution for semantic image segmentation. Proceedings of the 15th European Conference on Computer Vision. Munich, Germany: Springer, 2018,

- 833-851.
18. Vaswani, A. et al. Attention is all you need. Proceedings of the 31st International Conference on Neural Information Processing Systems. Long Beach, CA, USA: Curran Associates Inc., 2017, 6000-6010.
 19. Wang, X. L. et al. Non-local neural networks. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Salt Lake City, UT, USA: IEEE, 2018, 7794-7803.
 20. Zhu, P. F. et al. Vision meets drones: a challenge. Preprint at <https://doi.org/10.48550/arXiv.1804.07437> (2018).
 21. Lin, T. Y. et al. Microsoft COCO: common objects in context. Proceedings of the 13th European Conference on Computer Vision. Zurich, Switzerland: Springer, 2014, 740-755.
 22. Pang, J. M. et al. Libra R-CNN: towards balanced learning for object detection. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Long Beach, CA, USA: IEEE, 2019, 821-830.
 23. Ge, Z. et al. YOLOX: exceeding YOLO series in 2021. Preprint at <https://doi.org/10.48550/arXiv.2107.08430> (2021).
 24. Jocher, G. ultralytics/yolov5: v3.1 – Bug Fixes and Performance Improvements. (2020). At <https://github.com/ultralytics/yolov5> URL.
 25. Wang, C. Y., Bochkovskiy, A. & Liao, H. Y. M. YOLOv7: trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. Preprint at <https://doi.org/10.48550/arXiv.2207.02696> (2022).
 26. Feng, C. J. et al. TOOD: task-aligned one-stage object detection. Proceedings of the IEEE/CVF International Conference on Computer Vision. Montreal, QC, Canada: IEEE, 2021, 3490-3499.
 27. Ren, S. Q. et al. Faster R-CNN: towards real-time object detection with region proposal networks. Proceedings of the 28th International Conference on Neural Information Processing Systems. Montreal, Quebec, Canada: MIT Press, 2015, 91-99.
 28. Wang, J. Q. et al. CARAFE: content-aware reassembly of features. Proceedings of the IEEE/CVF International Conference on Computer Vision. Seoul, Korea (South): IEEE, 2019, 3007-3016.
 29. Lu, X. et al. Grid R-CNN. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Long Beach, CA, USA: IEEE, 2019, 7355-7364.
 30. Huang, Y. C., Chen, J. X. & Huang, D. UFMP-Det: toward accurate and efficient object detection on drone imagery. Preprint at <https://doi.org/10.48550/arXiv.2112.10415> (2021).
 31. Yang, F. et al. Clustered object detection in aerial images. Proceedings of the IEEE/CVF International Conference on Computer Vision. Seoul, Korea (South): IEEE, 2019, 8311-8320.
 32. Li, C. L. et al. Density map guided object detection in aerial images. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. Seattle, WA, USA: IEEE, 2020, 190-191.
 33. Deng, S. T. et al. A global-local self-adaptive network for drone-view object detection. *IEEE Transactions on Image Processing* **30**, 1556-1569 (2020).
 34. Wei, Z. W. et al. AMRNet: chips augmentation in aerial images object detection. Preprint at <https://doi.org/10.48550/arXiv.2009.07168> (2020).
 35. Rossi, L., Karimi, A. & Prati, A. A novel region of interest extraction layer for instance segmentation. Proceedings of the 25th International Conference on Pattern Recognition. Milan, Italy: IEEE, 2021, 2203-2209.
 36. Chen, Q. et al. You only look one-level feature. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Nashville, TN, USA: IEEE, 2021, 13039-13048.
 37. Chen, K. et al. MMDetection: open MMLab detection toolbox and benchmark. Preprint at <https://doi.org/10.48550/arXiv.1906.07155> (2019).
 38. Wang, J. Q. et al. Side-aware boundary localization for more precise object detection. Proceedings of the 16th European Conference on Computer Vision. Glasgow, UK: Springer, 2020, 403-419.
 39. Lin, T. Y. et al. Focal loss for dense object detection. Proceedings of the IEEE International Conference on Computer Vision. Venice, Italy: IEEE, 2017, 2980-2988.
 40. Xu, H. Y. et al. Deep regionlets for object detection. Proceedings of the 15th European Conference on Computer Vision. Munich, Germany: Springer, 2018, 827-844.
 41. Shrivastava, A. et al. Beyond skip connections: top-down modulation for object detection. Preprint at <https://doi.org/10.48550/arXiv.1612.06851> (2016).
 42. Zhang, S. F. et al. Single-shot refinement neural network for object detection. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Salt Lake City, UT, USA: IEEE, 2018, 4203-4212.
 43. Zhao, Q. J. et al. M2Det: a single-shot object detector based on multi-level feature pyramid network. Proceedings of the 33rd AAAI Conference on Artificial Intelligence. Honolulu, Hawaii, USA: AAAI, 2019, 9259-9266.