· Article · Advances in Polar Science
doi: 10.12429/j.advps.2025.0042 March 2026 Vol. 37 No. 1: 70-80

Retrieving rare aurora forms from all-sky images via synthetic-to-real progressive learning

ZHAI Chaoqiang & WANG Qian*

School of Communication and Information Engineering, Xi'an University of Posts and Telecommunications, Xi'an 710121, China

Received 4 November 2025; accepted 22 February 2026; published online 30 March 2026

Abstract Fine-scale structures can be observed in small field-of-view (FOV) auroral observations, but they are often overlooked because they appear only sporadically in all-sky observations. Such forms are of great interest because they may embody specific magnetosphere-ionosphere coupling processes, reveal localized energy deposition pathways, and provide new insights into cross-scale plasma dynamics and instabilities. However, their limited spatial extent, transient occurrence, and scarcity in wide-FOV observations make systematic investigation challenging. Traditional manual analysis struggles to capture these subtle structures within vast all-sky datasets, while automated detection faces severe data imbalance and morphological ambiguity. To address these challenges, we propose a synthetic-to-real progressive learning framework for cross-FOV retrieval of rare auroral forms. A Generative Adversarial Network (GAN) is employed to perform cross-FOV transformation between unpaired small-FOV images containing rare aurora forms and all-sky images (ASI) without such structures, thereby generating large numbers of synthetic ASI with rare auroral morphology. These synthetic samples are used to train an initial detection model, which subsequently undergoes iterative fine-tuning through feedback-guided learning: The model performs inference on new all-sky data, and the progressively accumulated real detections are incorporated into the training set. Experimental results demonstrate that the proposed method achieves over 92% detection accuracy on ASI, enabling high-precision retrieval of small-scale auroral structures across large-scale observations. This framework provides a scalable and effective approach to rediscovering rare auroral phenomena in continuous all-sky monitoring, offering new opportunities for exploring the fine-scale dynamics of the upper atmosphere.

Keywords **fine-scale auroral structures**, **rare auroral forms**, **cross-FOV retrieval**, **Generative Adversarial Network (GAN)**, **synthetic-to-real progressive learning**, **feedback-guided learning**

Citation: Zhai C Q, Wang Q. Retrieving rare aurora forms from all-sky images via synthetic-to-real progressive learning. Adv Polar Sci, 2026, 37(1): 70-80, doi: 10.12429/j.advps.2025.0042

1 Introduction

Auroras are natural manifestations of space weather and offer a direct window into the complex interactions between the solar wind (Shi et al., 2013; Tsurutani et al., 2023; Zirker, 1977) and the Earth's magnetosphere-ionosphere system (Buchert et al., 2008; Case et al., 2020; Liu et al., 2023; Nichols, 2011). The information carried by small-field-of-view (small-FOV) auroral observations, particularly fine-scale auroral structures and other rare auroral forms, is of significant scientific value, helping to reveal mechanisms of energy deposition and particle acceleration in the upper atmosphere (Ivarsen et al., 2024; Kataoka et al., 2021; Zhao et al., 2019). Although such rare forms appear sporadically in all-sky observations, their morphology is remarkably stable and distinct, indicating genuine physical processes rather than incidental artifacts. Investigating these features offers insights into fine-scale

* Corresponding author. ORCID: 0000-0002-3819-7369. E-mail: wqabby@xupt.edu.cn

plasma instabilities (Makarevich, 2016), localized energy-transfer pathways, and cross-scale coupling in magnetosphere-ionosphere dynamics, as well as connections to ionospheric disturbances relevant to space weather (Li et al., 2022). Recent work reports fragmented aurora-like emissions (FAEs), small, green, ripple-like structures near auroral arcs, likely linked to gradient-drift instability in the auroral ionosphere (Li et al., 2025). Figure 1 shows representative auroral imagery, including auroral photograph, small-FOV auroral image, and all-sky image (ASI).

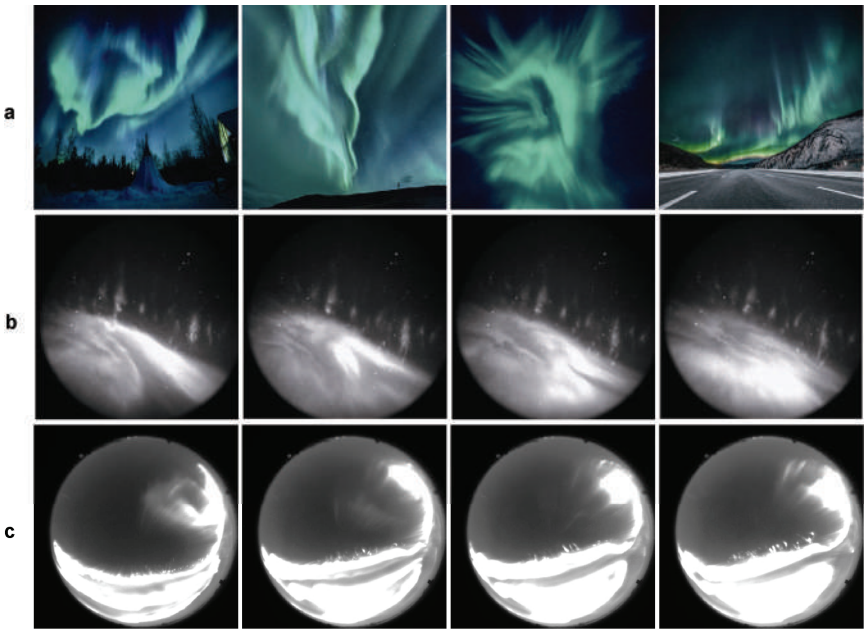

Figure 1 Auroral imagery. **a**, auroral photograph; **b**, small-FOV auroral image; **c**, all-sky image (ASI) containing fine-scale auroral structures.

Accurately detecting and analyzing these rare small-scale auroral structures remains a longstanding challenge. Their low occurrence frequency, limited spatial extent, and complex visual presentation render manual identification inefficient even for experienced researchers. Automated recognition faces similar difficulties due to severe data imbalance, dynamic backgrounds, and variable morphology, leaving such events underrepresented in large-scale morphological studies.

Small-FOV imaging provides high spatial resolution and superior signal-to-noise ratio for precise observation of detailed structures, yet its restricted spatiotemporal coverage hinders comprehensive monitoring. In contrast, all-sky image (ASI) provide continuous, wide-field coverage but often obscure fine structures amid complex background emissions. Bridging these complementary advantages motivates a cross-FOV retrieval strategy that integrates local precision with global continuity.

To address these challenges, we propose a framework for detecting and retrieving rare auroral forms in ASI. The approach leverages a Generative Adversarial Network (GAN) (Creswell et al., 2018) to transfer the morphological characteristics of rare auroral structures observed in small-FOV images into all-sky observations, generating synthetic all-sky data that contain these otherwise underrepresented features. A progressive learning scheme further refines the detection model: An object detector initially trained on small-FOV images is applied to the synthetic and new all-sky data, and the model is iteratively updated using progressively accumulated detections. This strategy alleviates the challenges of data scarcity and imbalance, improves generalization, and enables reliable retrieval of rare auroral events. While rare auroral forms are used here as a representative example, the framework is broadly applicable to other sparsely observed space physics phenomena. The main contributions of this work are as follows:

(1) We develop a novel model specifically designed for detecting and retrieving rare auroral structures in all-sky imagery, enabling systematic study of these fine-scale phenomena across large-scale datasets.

(2) We address the inherent contradiction between the rarity of certain auroral phenomena and the large data requirements of AI methods. By combining synthetic cross-FOV image generation with progressive real-sample adaptation, our approach mitigates data scarcity and imbalance, allowing effective model training without extensive manual labeling.

(3) We provide a generalizable technique for transferring knowledge from small-FOV observations to all-sky imagery. This cross-FOV methodology offers a practical tool for studying rare space physics events, facilitating broader applications in magnetosphere-ionosphere coupling research.

2 Related work

Automated classification and detection of auroral images, such as an automated auroral detection system (Nanjo et al., 2022; Wang et al., 2024) and performance optimization for auroral classification and detection (Wang Q et al., 2025), have been active research topics in space physics and atmospheric science. Existing algorithms primarily target common auroral forms such as arcs (Imajo et al., 2021) and curtains, but they perform poorly on rare small-FOV auroral structures, including fine-scale auroral structures such as fragmented aurora-like emissions (FAEs) (Dreyer et al., 2021; Li et al., 2025; Whiter et al., 2021) and spiral patterns (Huang et al., 2022). This limitation arises mainly from two factors. First, rare auroral events occur infrequently and persist for only short durations, resulting in an extremely limited number of labeled samples. Second, these structures are often confined to narrow observation regions and exhibit low contrast with blurred boundaries, making them easily obscured under limited resolution and signal-to-noise conditions. Traditional methods, such as threshold segmentation (Chen et al., 2025; Hu et al., 2024), morphological analysis (Kim et al., 2025), and handcrafted feature extraction (Alsakar et al., 2024; Sid'El Moctar et al., 2024), struggle in such scenarios, underscoring the need for advanced, data-driven techniques to enable generalized statistical studies of rare auroral forms.

To bridge the disparity between small-FOV high-resolution imagery and ASI with wide-field but lower per-pixel resolution, this study introduces a GAN-based cross-FOV transformation framework that transfers the morphological characteristics of small-FOV auroral structures from the local observation domain to the global domain. The generative network, while maintaining radiometric consistency and geometric fidelity, embeds fine spatial textures from small-FOV images into the visual-semantic space of ASI, thereby synthesizing realistic all-sky samples containing rare auroral forms such as FAEs. This process substantially expands the training set for underrepresented categories and produces pseudo-labels for downstream detectors, effectively alleviating data imbalance and mitigating distribution shifts between small-FOV and all-sky domains.

At the detection level, a feedback-guided learning strategy within a synthetic-to-real progressive learning scheme enables rapid model adaptation to new auroral forms using minimal labeled data. By continuously generating synthetic samples through GAN-based augmentation, the system introduces hard examples that improve robustness to low contrast, indistinct boundaries, and complex background textures. To further address cross-scale inconsistencies, multi-scale feature pyramids (Li et al., 2024; van Quyen and Kim, 2023; Yu et al., 2024) and domain adaptation mechanisms (Qian et al., 2025) are incorporated into the detector, improving cross-scale robustness and reducing performance degradation caused by variations in imaging geometry and spectral characteristics.

By integrating these strategies, the proposed hybrid framework enables efficient and stable cross-FOV retrieval of rare small-FOV auroral events from extensive all-sky datasets. This approach resolves fundamental challenges in rare aurora detection and establishes a methodological foundation for large-scale statistical analyses of magnetosphere-ionosphere coupling in high-latitude regions. By effectively bridging small-FOV high-resolution data with continuous all-sky observations, the method enhances our capacity to study the dynamic interactions governing auroral activity, offering new insights into energy transfer, particle acceleration, and space weather effects. While FAEs serve as the representative case, the framework is broadly applicable to other rare atmospheric or ionospheric events that are difficult to detect or image (e.g., mesospheric clouds and transient ionospheric disturbances (Collinson et al., 2019)), paving the way for more accurate, robust, and large-scale analyses of upper-atmospheric dynamics.

3 Methodology

3.1 Overview

This study aims to seamlessly transfer the recognition capability for rare auroral forms—previously validated in small-FOV imaging systems—to ASI, thereby enabling efficient and reliable detection of fine-scale auroral structures and other sub-kilometer features in all-sky imagery. Achieving this goal broadens observational coverage and establishes a data foundation for statistical studies of rare events and for analyzing magnetosphere-ionosphere coupling mechanisms. Two core challenges must be addressed.

(1) One is sample scarcity. Rare auroral events occur infrequently and last only briefly. Although more readily captured in small-FOV observations, they remain difficult to detect within large ASI datasets, accounting for <0.5% of publicly available or self-collected imagery. Moreover, existing ASI data seldom include finely labeled instances, making it infeasible to train conventional supervised frameworks directly.

(2) The other is cross-FOV and cross-scale differences. Substantial discrepancies between small-FOV and ASI-FOV size, spatial resolution, projection distortion, and signal-to-noise ratio—compress rare structures to only a few pixels and blend them with diffuse backgrounds, hindering both visual and algorithmic discrimination.

To overcome these challenges, as illustrated in Figure 2, we propose a "Generate-Adapt-Iterate" hybrid transfer framework. A bidirectional GAN for cross-FOV transformation (small-FOV ↔ ASI) performs radiometric and geometric consistency mappings on small-FOV auroral images, generating numerous pseudo-ASI samples that preserve the morphological and textural characteristics of fine-scale structures.

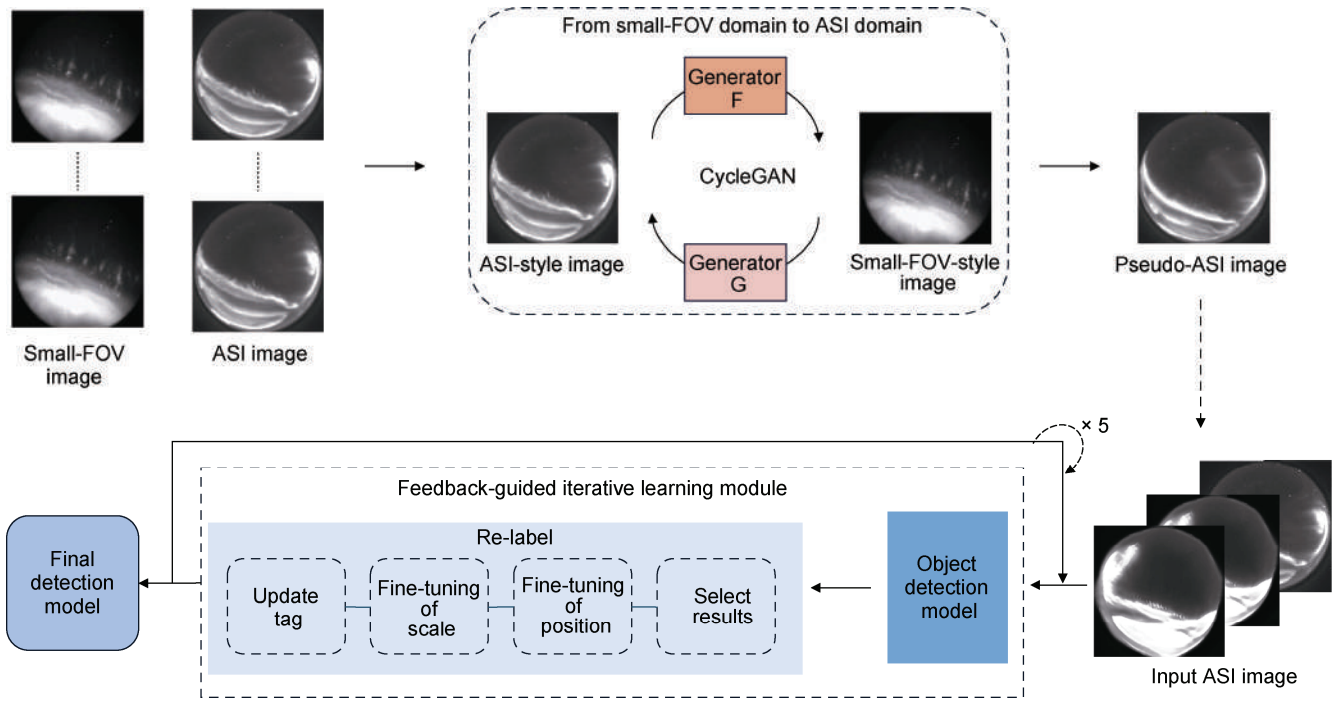

Figure 2 "Generate-Adapt-Iterate" hybrid transfer framework.

A pre-trained detection network (e.g., YOLOv9; Wang C Y et al., 2025) is first trained on labeled small-FOV images to learn the fine-scale features of rare auroral forms. This model provides an initial prior for subsequent transfer-adaptation on ASI data within a synthetic-to-real progressive learning scheme. During multi-round feedback-guided adaptation, the detector is applied to unlabeled real ASI, and pseudo-labels are selected by confidence thresholds. Only ~5% of difficult cases require manual correction, which are then used to curate a high-quality training set for the next iteration. After five adaptation rounds, detection performance improves from 17.4% to 58.7%.

During iterative optimization, self-distillation and multi-scale feature alignment are introduced after each adaptation cycle to mitigate cross-scale mismatches. In addition, cosine-annealing learning-rate scheduling and mixed-precision training accelerate convergence.

This framework is both data-efficient and scalable. The GAN module and the detection backbone can be replaced by more advanced generative or recognition architectures as they emerge. Moreover, the progressive-adaptation mechanism can be readily integrated with active learning or uncertainty sampling to further reduce manual annotation cost.

We primarily evaluate on a small-FOV dataset of FAEs; the proposed method remains generally applicable to other datasets of rare auroral forms.

3.2 Synthetic sample generation with CycleGAN

A GAN in the form of CycleGAN (Zhu et al., 2017) is employed to learn cross-FOV transformation from small-FOV auroral images to ASI using unpaired data. The core principle is to reduce distribution discrepancy between the source domain (small-FOV images containing FAEs) and the target domain (all-sky domain) via unsupervised learning. In doing so, the model learns radiometric and geometric mappings that preserve the morphology of fine-scale auroral structures while embedding them into the global context of all-sky observations. As illustrated in Figure 3, unpaired datasets from both domains are used to synthesize realistic ASI samples that faithfully represent small-FOV structures within a wide-field view.

CycleGAN comprises two generators that establish a bidirectional mapping between the domains. Generator G translates small-FOV images to ASI, while Generator F performs the inverse transformation from ASI back to small-FOV. Through this dual mapping (along with the associated adversarial objectives and cycle constraints), the framework promotes radiometric and geometric consistency in the synthesized images, preserving both visual appearance and spatial structure critical for downstream detection.

During generation, maintaining style coherence and spatial fidelity is essential to ensure that synthetic ASI are both realistic and compatible with the detector trained under the synthetic-to-real progressive learning scheme. To this end, we adopt the following strategies.

(1) Style alignment (adversarial objective). To maintain consistent visual characteristics—texture, brightness, and color distribution—between synthesized and real ASI while retaining the morphology of fine-scale auroral structures (e.g., FAEs), we apply standard adversarial losses.

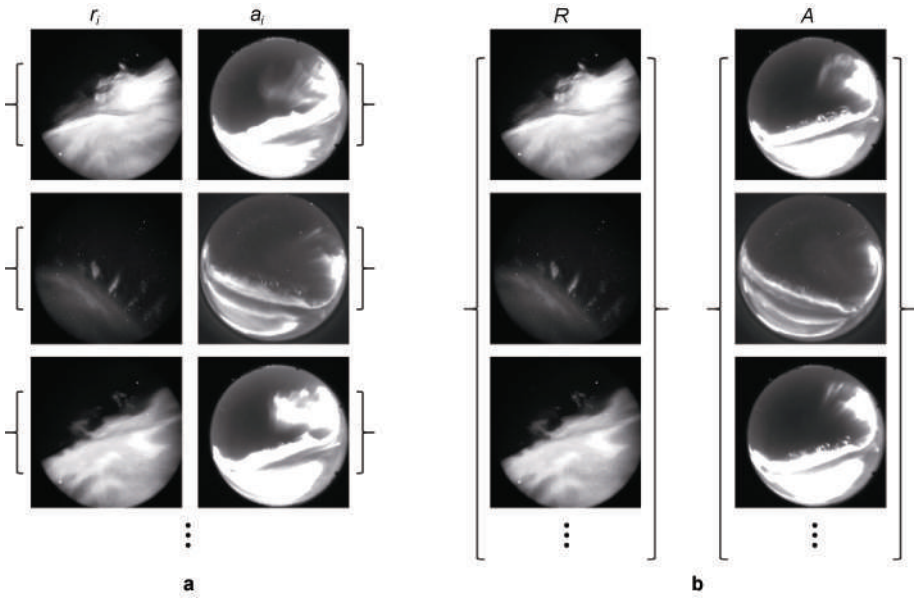

Figure 3 Illustration of data pairing strategies between small-FOV and all-sky auroral images. **a**, paired small-FOV and all-sky data; **b**, unpaired small-FOV and all-sky data.

$$L_{\mathrm{GAN}}\left(G, D_A, X, Y\right) = E_{y\sim p_{\mathrm{data}}(y)}\left[\log D_A(y)\right] + E_{x\sim p_{\mathrm{data}}(x)}\left[\log\left(1 - D_A\left(G(x)\right)\right)\right], \tag{1}$$

$$L_{\mathrm{GAN}}\left(F, D_R, Y, X\right) = E_{x\sim p_{\mathrm{data}}(x)}\left[\log D_R(x)\right] + E_{y\sim p_{\mathrm{data}}(y)}\left[\log\left(1 - D_R\left(F(y)\right)\right)\right], \tag{2}$$

$$L_{\mathrm{GAN}}(G, F) = L_{\mathrm{GAN}}\left(G, D_A, X, Y\right) + L_{\mathrm{GAN}}\left(F, D_R, Y, X\right), \tag{3}$$

where, $X = \{x\}$ denotes the set of small-FOV auroral images (containing FAEs), and $Y = \{y\}$ denotes the set of ASI. Generator G learns the mapping from the small-FOV domain R to the all-sky domain A, while generator F performs the inverse mapping $A \rightarrow R$. Discriminator D_A determines whether an input belongs to the real all-sky domain, and D_R evaluates authenticity in the small-FOV domain. The expectations $E_{y\sim p_{\mathrm{data}}(y)}$ and $E_{x\sim p_{\mathrm{data}}(y)}$ are taken over the true data distributions of the ASI and small-FOV domains, respectively, from the true data distribution of the all-sky domain; $E_{x\sim p_{\mathrm{data}}(y)}$ represents the expectation calculated by sampling real data x from the true data distribution of the small-FOV image domain.

(2) The cycle consistency loss is employed to preserve the spatial structure of the generated images during cross-FOV transformation. Specifically, an ASI produced by generator G, when mapped back by generator F into a small-FOV image, should remain consistent with the original small-FOV input; conversely, a small-FOV image translated by F and then reconstructed by G should recover the original ASI observation. This bidirectional constraint enforces spatial and structural coherence across the translation process, enabling both generators to maintain geometric integrity and the morphology of fine-scale auroral structures (e.g., FAEs).

$$L_{\mathrm{cyc}}(G, F) = E_{x\sim p_{\mathrm{data}}(x)}\left[\left\| F(G(x)) - x\right\|_1\right] + E_{y\sim p_{\mathrm{data}}(y)}\left[\left\| G(F(y)) - y\right\|_1\right], \tag{4}$$

$$L_{\mathrm{Total}}(G, F, D_R, D_A) = L_{\mathrm{GAN}}\left(G, D_A, X, Y\right) + L_{\mathrm{GAN}}\left(F, D_R, Y, X\right) + \lambda_{\mathrm{cyc}} L_{\mathrm{cyc}}(G, F), \tag{5}$$

where, $\left\| \ \right\|_1$ denotes the L1 norm; $X = \{x\}$ and $Y = \{y\}$ are the sets of small-FOV and ASI images, with expectations taken over the true data distributions $P_{\mathrm{data}}(x)$ and $P_{\mathrm{data}}(y)$, respectively. The weight λ_{cyc} balances adversarial realism and structural preservation and is set to 10. L_{GAN} is the adversarial loss, and L_{cyc} is the cycle consistency loss.

These strategies jointly ensure the visual realism of the synthesized images while preserving spatial structure, allowing the generated ASI samples to be used as effective training data for subsequent detection.

Given the potential photometric disparities between small-FOV and ASI, we adopt a brightness control (radiometric normalization) module. This procedure normalizes the brightness and contrast distribution to reduce training bias arising from photometric mismatches. As a result, the brightness characteristics of synthesized images are aligned with those of real ASI, thereby enhancing the detector's adaptability and robustness to photometric variability.

As shown in Figure 4, the synthesized ASI samples contain rich FAE fine-scale auroral structures that closely

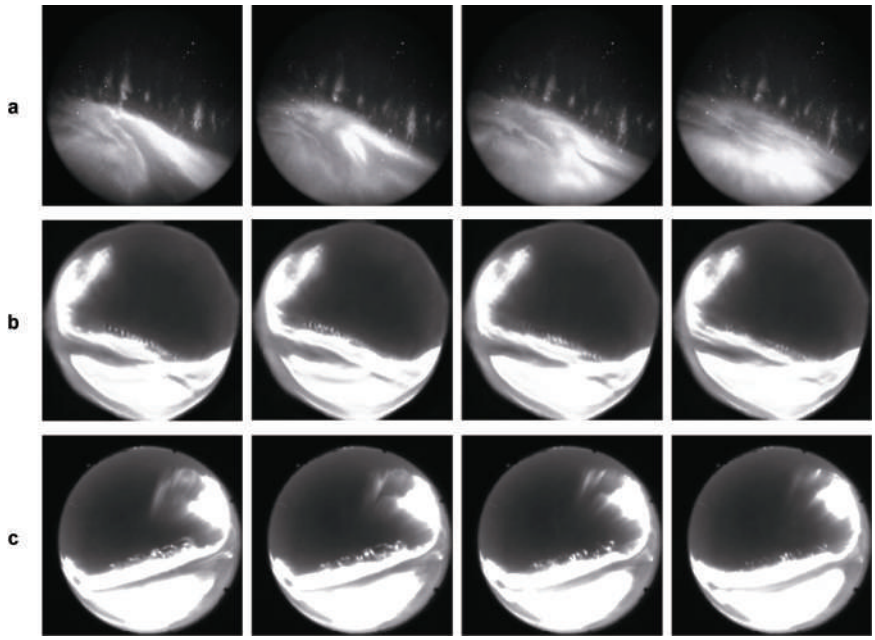

Figure 4 Comparison between small-FOV observations, synthesized ASI, and real ASI. **a**, small-FOV image containing FAEs fine-scale auroral structures; **b**, synthesized ASI from the cross-FOV transformation; **c**, real ASI.

resemble real ASI in both form and content. These synthetic images substantially expand the originally scarce positive samples for rare auroral forms, providing sufficient data for training the downstream object detection model.

3.3 Initial detector training

During the initial training stage, small-FOV images containing FAEs are used to pretrain the object-detection model so that it learns discriminative morphology for rare auroral forms. A large number of ASI samples containing FAEs are then synthesized via a CycleGAN-based cross-FOV transformation, providing rich training data for detection. The resulting dataset comprises two sources: (1) synthetic ASI images produced by CycleGAN that embed the characteristic small-FOV morphology, and (2) annotated real ASI. Although the number of truly rare events remains limited, the real images still supply valuable foundational supervision.

To enhance robustness and generalization, we apply standard data augmentation (rotation, flipping, scaling) to further expand the training set. We adopt efficient object-detection architectures—such as the YOLO (Hussain, 2024; Redmon et al., 2016) family and RT-DETR (Wang S et al., 2025; Zhao et al., 2024)—as the basis for the preliminary detector; in this study, YOLOv9 is used as the initial model.

The pretrained small-FOV detector is subsequently applied to the constructed ASI dataset for preliminary inference. Having learned the distinctive features of FAEs structures from small-FOV data, the model predicts bounding boxes and confidence scores on ASI; these serve as pseudo-labels for downstream training within a synthetic-to-real progressive learning pipeline.

Despite combining synthetic and real supervision, the preliminary evaluation exposes a notable domain-shift performance gap when transferring from small-FOV images to ASI, with especially poor recall on rare FAE structures. To address this, we introduce a feedback-guided, multi-round fine-tuning strategy that progressively adapts the detector to ASI, incrementally improving cross-FOV retrieval performance.

3.4 Progressive adaptation with fine-tuning

To further improve generalization and accuracy, we adopt a feedback-guided, synthetic-to-real progressive learning strategy (Figure 5). The goal is to iteratively optimize the detector through multiple rounds of adaptation-fine-tuning, thereby enhancing cross-FOV retrieval of rare auroral forms in ASI.

(1) Progressive adaptation. A pretrained detector (initialized on small-FOV images with FAEs fine-scale auroral structures) is first deployed over the constructed ASI dataset to generate bounding boxes and confidence scores. Although this initial model cannot perfectly identify all targets in the all-sky domain, it provides valuable localization and class priors. We then select high-confidence detections via a threshold to filter out low-confidence results; the retained boxes are converted into pseudo-labels that guide subsequent training and help the model further learn the characteristics of rare structures.

(2) Fine-tuning with limited supervision. To improve pseudo-label quality, approximately 5% of the labels are manually corrected to refine box position and scale, ensuring accurate supervision. The corrected pseudo-labels,

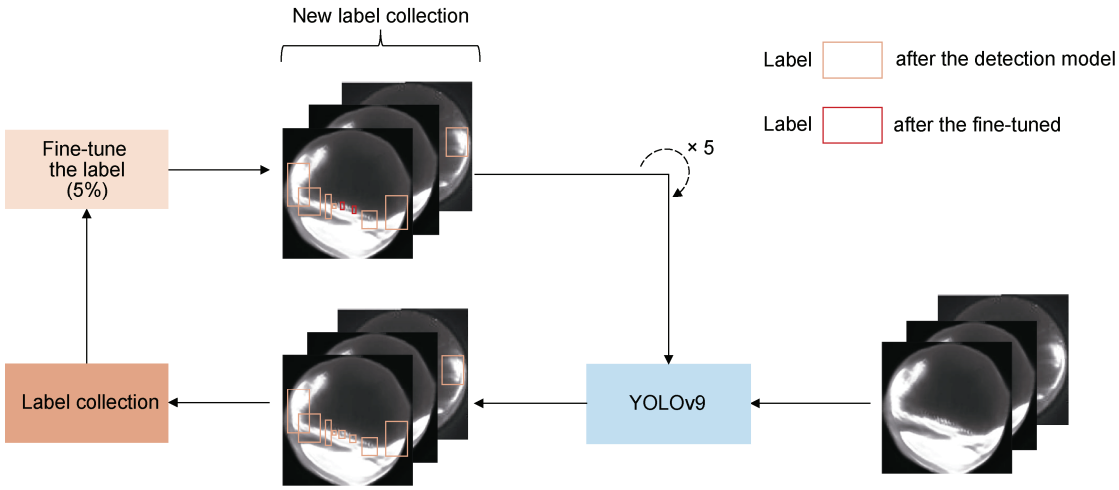

Figure 5 Progressive learning driven by iterative detection in ASI.

together with the remaining labels, are used to retrain the detector. We perform multi-round fine-tuning: After each round, the updated model re-infers on real ASI images, pseudo-labels are refreshed, and training proceeds to the next cycle. Through these iterations, detection performance gradually improves and the model adapts to diverse ASI conditions.

(3) Optimization schedule. A dynamically adjusted learning rate (e.g., cosine-annealing) is employed to accelerate convergence while mitigating overfitting: as training epochs increase, the learning rate smoothly decays, enabling fine-grained adjustments near the optimum.

(4) Evaluation. After each round, we assess performance using standard detection metrics: mean Average Precision (mAP), precision, and recall. Experiments show that, as fine-tuning progresses, the detector's accuracy on fine-scale auroral structures consistently improves, enabling accurate identification and localization of rare auroral forms across a broader range of ASI.

4 Experimental analysis

This section presents an in-depth evaluation of the proposed experimental strategy and model, together with ablation studies, to enable a comprehensive assessment of the method. We examine the cross-FOV retrieval framework—including the GAN-based synthesis pipeline and the feedback-guided, synthetic-to-real progressive learning scheme—and analyze results using standard detection metrics.

4.1 Dataset

The experimental data were acquired at the Antarctic Zhongshan Station. The dataset comprises ASI and small-FOV observations collected by the Polar Research Institute of China (PRIC) at the G-band (OI 557.7 nm). Continuous observations over 23 years (1995–2018) provide a substantial empirical basis for this study. Each raw image has a resolution of 1,024×1,024 pixels. Within the CycleGAN pipeline, inputs are standardized to 256×256 pixels by first resizing to 286×286 pixels and then cropping to 256×256 pixels. The cross-FOV synthesis is trained on two unpaired datasets, each comprising 1,000 images (1,000 small-FOV and 1,000 ASI).

For pretraining the detector, we use 2,886 small-FOV images containing FAEs: 2,308 for training and 578 for testing. In addition, the detection training corpus includes 424 ASI images with clear fine-scale auroral structures (324 train and 100 test). To increase diversity, data augmentation (rotation and brightness adjustment) is applied during training.

Data augmentation improves the model's capacity to learn invariant features and thereby enhances generalization to unseen data. Given the limited number of exemplars with characteristic morphology in auroral datasets, augmentation is essential to mitigate overfitting and improve robustness. We find that applying horizontal flipping, rotation, and brightness variation markedly boosts detection of FAEs in ASI, which is particularly beneficial for rare auroral forms that appear as small targets within a wide field of view.

4.2 Evaluation metrics

The evaluation metrics used in this study include precision, recall, and mean Average Precision as the primary indicators of detection performance. Precision measures the proportion of predicted positives that are correct, while recall measures the proportion of ground-truth positives that are successfully detected.

$$P = \frac{T}{T+F} \tag{6}$$

$$R = \frac{T}{T+N} \tag{7}$$

where, T (true positives) are correct detections whose predicted bounding boxes match the ground truth; F (false

positives) are incorrect detections; and N (false negatives) are missed targets.

$$m = \frac{1}{N}\sum_{i=1}^{N} A_i \tag{8}$$

where m (mAP) evaluates overall performance by averaging A (Average Precision) across categories. We report two standard variants: m_{50} (mAP_{50}), computed at a single Intersection-over-Union (IoU) threshold of 0.5, and $m_{50:95}$ ($\mathrm{mAP}_{50:95}$), computed by averaging A over IoU thresholds from 0.50 to 0.95 with a step of 0.05.

4.3 Results

Experiments were implemented in Python with PyTorch 1.9.1 and executed on an NVIDIA GeForce RTX 3080 Ti GPU, ensuring efficient training and inference. We evaluate the proposed approach through a series of ablation and comparative studies. The results demonstrate effective cross-FOV retrieval and accurate detection of rare auroral forms such as FAEs in ASI.

4.3.1 Cross-FOV retrieval of rare fine-scale auroral structures in ASI

To validate the effectiveness and stability of the proposed method on ASI, we conducted detection experiments on a continuous time sequence, as shown in Figure 6. Figure 6 provides a time-series visualization of the detection results, illustrating the Cross-FOV retrieval of rare fine-scale auroral structures across consecutive frames and explicitly reflecting their temporal dynamics, including emergence, evolution, and dissipation. Despite the fact that FAEs in ASIs are typically small in scale, low in contrast, and easily blended with background emissions, our method maintains a relatively consistent detection response over the continuous time window, enabling persistent retrieval and localization of FAEs. These results indicate that the

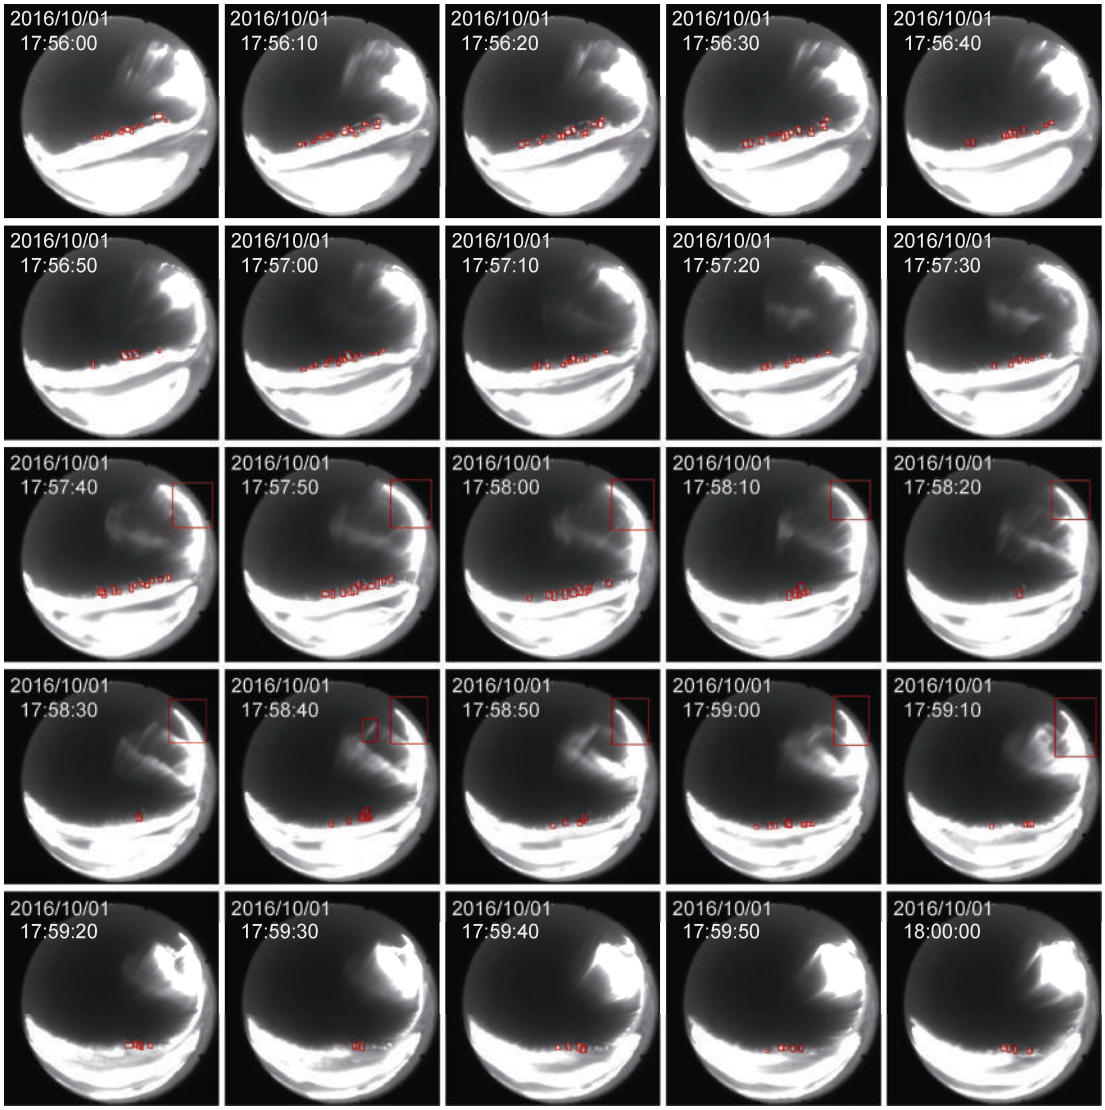

Figure 6 Representative retrieval results for FAEs in ASI. Red boxes indicate detected auroral structures that are likely FAEs.

proposed approach is well suited for fine-scale and rare auroral structure retrieval in complex time-varying scenarios, where such subtle features are often difficult to capture reliably using conventional imaging and standard detection pipelines.

4.3.2 Comparison with existing detection approaches

To validate the effectiveness of the proposed cross-FOV recognition framework for rare auroral forms, we compare it with representative auroral image-processing methods, including traditional feature-based approaches and advanced deep-learning detectors. Traditional pipelines exhibit inherent limitations in feature extraction, while current deep-learning models (e.g., YOLO, RT-DETR) still face challenges when encountering complex auroral morphology.

Our method integrates a GAN-based cross-FOV transformation (CycleGAN) with an object-detection framework, specifically optimized for FAEs. By introducing synthetic image generation and feedback-guided progressive adaptation, the approach performs well on synthetic data, and improves crucially generalization on real ASI. Comparative experiments show that traditional methods (e.g., thresholding, morphological processing) yield low accuracy on small targets, whereas a baseline YOLO detector performs well in simple scenes but is susceptible to background noise, resolution limits, and deformation in complex auroral imagery.

Concretely, existing methods exhibit the following shortcomings: YOLO detects targets reliably in simple images; however, accuracy degrades markedly for rare FAEs under strong noise or shape distortion.

RT-DETR enhances flexibility with a Transformer backbone, yet recall remains insufficient on ASI due to large field-of-view disparity and complex backgrounds—limitations that are amplified in the presence of intricate FAEs.

By contrast, the proposed cross-FOV framework—through synthetic augmentation and progressive adaptation (Figure 7)—effectively addresses these issues, particularly under complex backgrounds and for rare structures, yielding substantial gains in both precision and recall. As shown in Table 1, the hybrid design that combines generative modeling with discriminative detection achieves strong performance on ASI containing FAEs, demonstrating clear potential for efficient and accurate processing of auroral imagery.

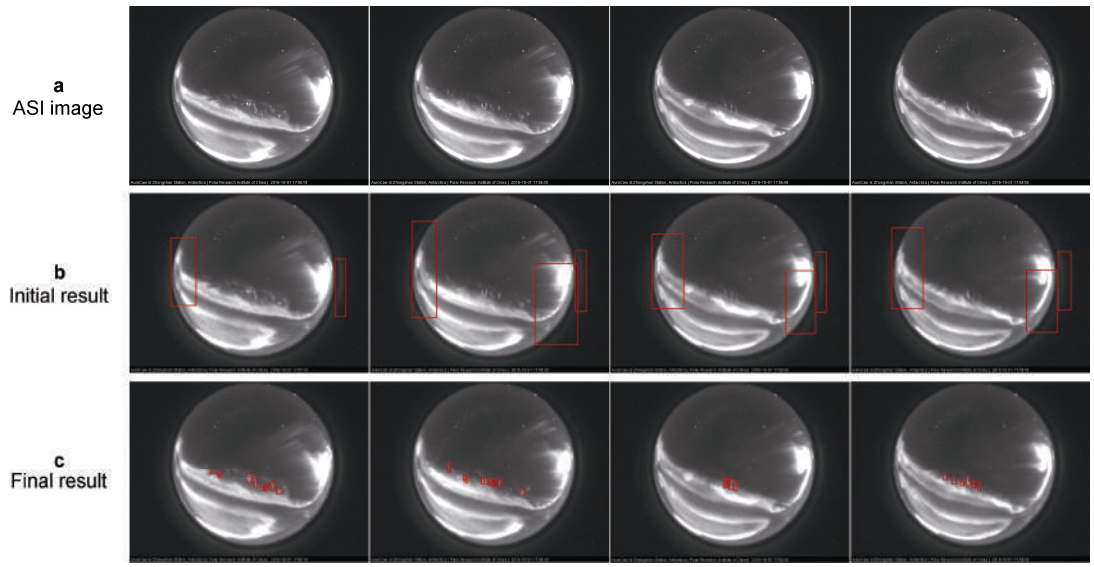

Figure 7 Comparison of detection results during synthetic-to-real progressive learning: **a**, original ASI image; **b**, inference of the pretrained small-FOV detector; **c**, inference after five rounds of feedback-guided fine-tuning. Red boxes indicate detected auroral structures that are likely FAEs.

Table 1 Detection performance of advanced object-detection models on ASI images containing FAEs structures

Model	Parameters/M	FLOPs/G	Precision/%	Recall/%	m_{50} value	$m_{50:95}$ value
RT-DETR-R34	31.0	92.0	94.0	95.1	95.3	56.7
YOLOv5-L r7.0	46.5	109.1	93.8	94.6	95.2	56.5
YOLOv7	36.9	104.7	92.7	91.5	94.9	50.6
YOLOv8-M	25.9	78.9	94.4	95.4	95.8	57.1
YOLOv9-C	25.3	102.1	94.8	96.0	96.0	57.8
YOLOv10-B	20.1	98.7	94.2	95.6	95.7	57.5

4.3.3 Component ablation and parameter sensitivity studies

To quantify the contribution of each component in the proposed cross-FOV retrieval framework for rare auroral forms, we conduct ablation experiments that progressively evaluate (1) the GAN-based cross-FOV synthesis module and (2) the feedback-guided progressive fine-tuning scheme. The comparative results (Table 2) show that both components are crucial: the synthesis module expands positive samples and reduces domain shift, while the feedback-guided adaptation further improves detection on ASI by refining pseudo-labels across rounds.

Table 2 Ablation of the GAN-based cross-FOV synthesis and feedback-guided progressive fine-tuning modules (all experiments use YOLOv9-C)

Cycle	Fine-tuning	Precision/%	Recall/%	$m_{50:95}$ value	$m_{50:95}$ test
		39.4	40.8	17.4	5.5
√		47.7	49.1	21.4	20.6
	√	95.5	96.4	58.7	31.7
√	√	95.5	96.4	58.7	57.9

We also examine parameter sensitivity with respect to the number of adaptation rounds and the use of generative fine-tuning (Table 3). As the number of iterations increases, the detector performance (measured by mAP, precision, and recall) improves steadily, with the best results at five rounds. Accordingly, we adopt five iterations as the default setting.

Table 3 Sensitivity of feedback-guided fine-tuning to the number of adaptation rounds (mAP, precision, and recall; all experiments use YOLOv9-C)

Iteration number	Precision/%	Recall/%	m_{50} value	$m_{50:95}$ value	$m_{50:95}$ test
0	39.4	40.8	42.1	17.4	16.6
1	70.7	69.2	69.9	36.3	36.0
2	84.8	85.4	86.1	50.5	49.8
3	90.2	89.1	89.7	56.4	56.0
4	94.8	94.7	95.5	58.1	57.6

5 Conclusions

This paper presents a cross-FOV recognition framework for rare auroral forms, improving the detection of FAEs in small-FOV imagery and transferring this capability to ASI. The method integrates a GAN (implemented with CycleGAN) and an object-detection pipeline, addressing sample scarcity and cross-FOV differences via synthetic sample augmentation and feedback-guided, synthetic-to-real progressive learning. Experiments demonstrate that combining generative augmentation with progressive adaptation substantially enhances performance on ASI: with mixed training on synthetic and real images, the detector achieves notable gains in both accuracy and recall for rare FAEs structures. While evaluations focus on FAEs cases, the framework is morphology-agnostic and readily applicable to other rare auroral forms, enabling scalable, high-precision cross-FOV retrieval in continuous all-sky monitoring.

Acknowledgments This work was supported by the National Natural Science Foundation of China (Grant no. 41874173).

References

Alsakar Y M, Sakr N A, Elmogy M. 2024. An enhanced classification system of various rice plant diseases based on multi-level handcrafted feature extraction technique. Sci Rep, 14(1): 30601, doi:10.1038/s41598-024-81143-1.

Buchert S C, Tsuda T, Fujii R, et al. 2008. The Pedersen current carried by electrons: a non-linear response of the ionosphere to magnetospheric forcing. Ann Geophys, 26(9): 2837-2844, doi:10.5194/angeo-26-2837-2008.

Case N A, Grocott A, Fear R C, et al. 2020. Convection in the magnetosphere-ionosphere system: a multimission survey of its response to IMF by reversals. J Geophys Res Space Phys, 125(10): e2019JA027541, doi:10.1029/2019JA027541.

Chen X, Liu C, Xie D, et al. 2025. Image thresholding segmentation method based on adaptive granulation and reciprocal rough entropy. Inf Sci, 695: 121737, doi:10.1016/j.ins.2024.121737.

Collinson G, McFadden J, Mitchell D, et al. 2019. Traveling ionospheric disturbances at Mars. Geophys Res Lett, 46(9): 4554-4563, doi:10.1029/2019GL082412.

Creswell A, White T, Dumoulin V, et al. 2018. Generative adversarial networks: an overview. IEEE Signal Process Mag, 35(1): 53-65, doi:10.1109/MSP.2017.2765202.

Dreyer J, Partamies N, Whiter D, et al. 2021. Characteristics of fragmented aurora-like emissions (FAEs) observed on Svalbard. Ann Geophys, 39(2): 277-288, doi:10.5194/angeo-39-277-2021.

Hu P, Han Y B, Zhang Z, et al. 2024. A multi-level thresholding image segmentation algorithm based on equilibrium optimizer. Sci Rep, 14: 29728, doi:10.1038/s41598-024-81075-w.

Huang K, Liu Y H, Lu Q M, et al. 2022. Auroral spiral structure formation through magnetic reconnection in the auroral acceleration region. Geophys Res Lett, 49(18): e2022GL100466, doi:10.1029/2022GL100466.

Hussain M. 2024. YOLOv1 to v8: unveiling each variant–a comprehensive review of YOLO. IEEE Access, 12: 42816-42833, doi:10.1109/ACCESS.2024.3378568.

Imajo S, Miyoshi Y, Kazama Y, et al. 2021. Active auroral arc powered by accelerated electrons from very high altitudes. Sci Rep, 11(1): 1610, doi:10.1038/s41598-020-79665-5.

Ivarsen M F, Huyghebaert D R, Gillies M D, et al. 2024. Turbulence around auroral arcs. J Geophys Res Space Phys, 129(8): e2023JA032309, doi:10.1029/2023JA032309.

Kataoka R, Chaston C C, Knudsen D, et al. 2021. Small-scale dynamic aurora. Space Sci Rev, 217(1): 17, doi:10.1007/s11214-021-00796-w.

Kim V, Adaloglou N, Osterland M, et al. 2025. Self-supervision advances

morphological profiling by unlocking powerful image representations. Sci Rep, 15(1): 4876, doi:10.1038/s41598-025-88825-4.

Li B, Zong Q G, Hu Z J, et al. 2025. Fragmented aurora-like emissions–a visualization of aurora aroused ionospheric instability ripples. J Geophys Res Space Phys, 130(11): e2025JA034225, doi:10.1029/2025JA034225.

Li W, Zhao D S, He C Y, et al. 2022. Spatial-temporal behaviors of large-scale ionospheric perturbations during severe geomagnetic storms on September 7–8 2017 using the GNSS, SWARM and TIE-GCM techniques. J Geophys Res Space Phys, 127(3): e2021JA029830, doi:10.1029/2021JA029830.

Li X B, Yu H F, Chen H Y. 2024. Multi-scale aggregation feature pyramid with cornerness for underwater object detection. Vis Comput, 40(2): 1299-1310, doi:10.1007/s00371-023-02849-3.

Liu J, Wang W B, Qian L Y, et al. 2023. Impacts of ionospheric conductance on magnetosphere-ionosphere coupling. J Geophys Res Space Phys, 128(2): e2022JA030864, doi:10.1029/2022JA030864.

Makarevich R A. 2016. Toward an integrated view of ionospheric plasma instabilities: Altitudinal transitions and strong gradient case. J Geophys Res Space Phys, 121(4): 3634-3647, doi:10.1002/2016JA022515.

Nanjo S, Nozawa S, Yamamoto M, et al. 2022. An automated auroral detection system using deep learning: real-time operation in Tromsø, Norway. Sci Rep, 12(1): 8038, doi:10.1038/s41598-022-11686-8.

Nichols J D. 2011. Magnetosphere–ionosphere coupling at Jupiter-like exoplanets with internal plasma sources: implications for detectability of auroral radio emissions. Mon Not R Astron Soc, 414(3): 2125-2138, doi:10.1111/j.1365-2966.2011.18528.x.

Qian Q, Luo J, Qin Y. 2025. Adaptive intermediate class-wise distribution alignment: a universal domain adaptation and generalization method for machine fault diagnosis. IEEE Trans Neural Netw Learn Syst, 36(3): 4296-4310, doi:10.1109/TNNLS.2024.3376449.

Redmon J, Divvala S, Girshick R, et al. 2016. You only look once: unified, real-time object detection. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Las Vegas: IEEE, 779-788, doi:10.1109/CVPR.2016.91.

Shi Q Q, Zong Q G, Fu S Y, et al. 2013. Solar wind entry into the high-latitude terrestrial magnetosphere during geomagnetically quiet times. Nat Commun, 4: 1466, doi:10.1038/ncomms2476.

Sid'El Moctar S M, Rida I, Boudaoud S. 2024. Comprehensive review of feature extraction techniques for sEMG signal classification: from handcrafted features to deep learning approaches. IRBM, 45(6): 100866, doi:10.1016/j.irbm.2024.100866.

Tsurutani B T, Zank G P, Sterken V J, et al. 2023. Space plasma physics: a review. IEEE Trans Plasma Sci, 51(7): 1595-1655, doi:10.1109/TPS.2022.3208906.

van Quyen T, Kim M Y. 2023. Feature pyramid network with multi-scale prediction fusion for real-time semantic segmentation. Neurocomputing, 519(C): 104-113, doi:10.1016/j.neucom.2022.11.062.

Wang C Y, Yeh I H, Mark Liao H Y. 2025. YOLOv9: Learning what you want to learn using programmable gradient information//Leonardis A, Ricci E, Roth S, et al. (eds). Computer Vision – ECCV 2024. Lecture Notes in Computer Science, vol 15089. Cham: Springer, doi:10.1007/978-3-031-72751-1_1.

Wang Q, Bai W Y, Zhang W, et al. 2024. Automatically sketching auroral skeleton structure in all-sky image for measuring aurora arcs. J Geophys Res Space Phys, 129(4): e2023JA031778, doi:10.1029/2023JA031778.

Wang Q, Jing S H, Yang R, et al. 2025. Loosen attention: Integrating localized channel and coarse spatial attention for enhanced analysis of complex aurora images. Eng Appl Artif Intell, 160: 111907, doi:10.1016/j.engappai.2025.111907.

Wang S, Xia C L, Lv F, et al. 2025. RT-DETRv3: real-time end-to-end object detection with hierarchical dense positive supervision. 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). Tucson: IEEE, 1628-1636, doi:10.1109/WACV61041.2025.00166.

Whiter D K, Sundberg H, Lanchester B S, et al. 2021. Fine-scale dynamics of fragmented aurora-like emissions. Ann Geophys, 39(6): 975-989, doi:10.5194/angeo-39-975-2021.

Yu Y, Zhang Y, Cheng Z Y, et al. 2024. Multi-scale spatial pyramid attention mechanism for image recognition: an effective approach. Eng Appl Artif Intell, 133: 108261, doi:10.1016/j.engappai.2024.108261.

Zhao H Y, Zhou X Z, Zong Q G, et al. 2019. Small-scale aurora associated with magnetospheric flow vortices after a solar wind dynamic pressure decrease. J Geophys Res Space Phys, 124(5): 3303-3311, doi:10.1029/2018ja026234.

Zhao Y A, Lv W Y, Xu S L, et al. 2024. DETRs beat YOLOs on real-time object detection. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Seattle: IEEE, 16965-16974, doi:10.1109/CVPR52733.2024.01605.

Zhu J Y, Park T, Isola P, et al. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. 2017 IEEE International Conference on Computer Vision (ICCV). Venice: IEEE, 2242-2251, doi:10.1109/ICCV.2017.244.

Zirker J B. 1977. Coronal holes and high-speed wind streams. Rev Geophys, 15(3): 257-269, doi:10.1029/RG015i003p00257.