Detecting Aurora in the Near-Infrared from Space: A Deep Learning Approach with Explainable AI (Space Weather under review)

Junmu Youn (jmyoun@khu.ac.kr)[1], Serin Jeon[2,3], Woong Jeon[1,3], Se-Heon Jeong[3], Jeong-Heon Kim[3], Woo Kyoung Lee[2,3], and Yong-Jae Moon[1], Geonhwa Jee[2,4]

[1] Kyung Hee University, Korea [2] University of Science and Technology, Korea [3] Korea Astronomy and Space Science Institute, Korea [4] Division of Ocean and Atmosphere Sciences, Korea Polar Research Institute, Korea

Key Points

1. First deep-learning auroral classification from e-POP/FAI near-infrared time-series.
2. Temporal gaps let the model learn auroral dynamics, filtering out clouds and city lights.
3. Explainable AI localizes auroral structures without pixel-level spatial supervision.

Introduction

Auroras form along the auroral oval where magnetospheric particles strike the upper atmosphere, and the boundary of the oval reflects geomagnetic activity and space weather.

Space weather disrupts radio communication and satellite navigation and can drive geomagnetically induced currents (GICs) (e.g., the May 2024 superstorm; Lawrence et al., 2025)

Auroral detection is essential because the aurora directly marks space weather that can disrupt communication, navigation, and power grids.

Goal This study aims to **automatically detect auroras** and estimate their location in satellite-based images to help determine the auroral oval boundary from its morphology.

Background Previous algorithms for auroral oval boundary determination have relied on ultraviolet (UV) imagers, such as TIMED/GUVI and DMSP/SSUSI, or visible imagers from ground-based all-sky cameras.

Challenge From **a single-channel, near infrared** satellite imager (e-POP/FAI), **it is difficult to distinguish faint auroral emissions from other bright features, such as clouds and city lights.**

Solution We **develop a deep learning model** that uses a sequence of images to **differentiate auroras from other features**, and an **Explainable AI (XAI) method** to visualize and interpret where the model detects the aurora.

Method

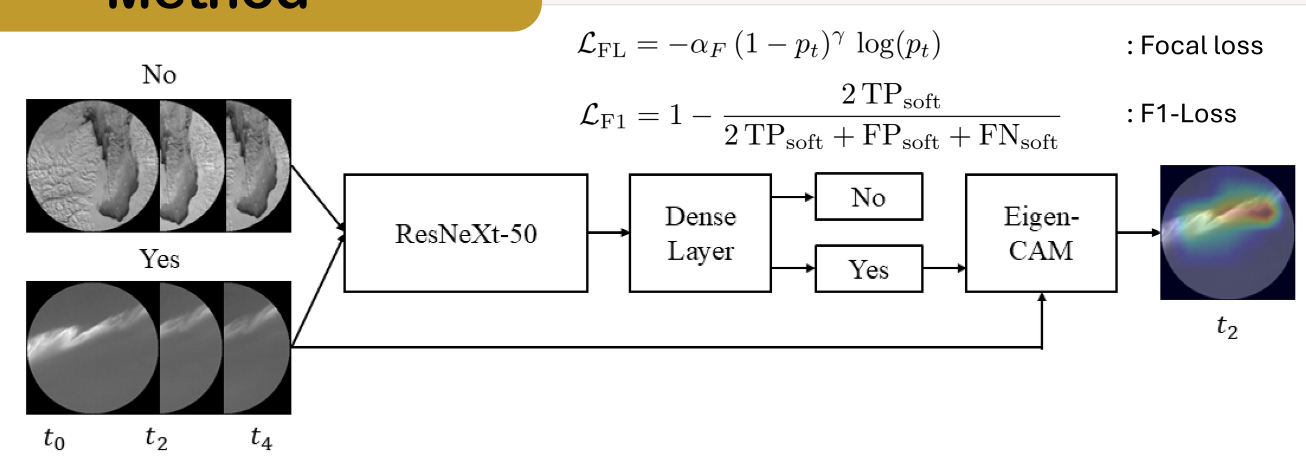

Figure 1. High-level flowchart of our study. The inputs are examples from our best-performing model

Deep learning model **ResNeXt-50** model (Xie et al., 2017), where the final fully connected layer is modified to output confidence scores for "Yes" and "No" classes.

Loss functions A combination of Focal Loss ($\mathcal{L}_{\mathrm{FL}}$) with label smoothing and F1-Loss ($\mathcal{L}_{\mathrm{F1}}$),

$$\boxed{\mathcal{L}_{\mathrm{total}} = \alpha \cdot \mathcal{L}_{\mathrm{FL}} + (1-\alpha)\cdot \mathcal{L}_{\mathrm{F1}}}$$

where $\alpha =$ 0.7. For the first seven epochs, only Focal Loss is applied.

Five input configurations compared (Models 1 to 5) vary frame count and cadence, namely **(1)** single image, **(2)** two images, **(3)** three images, **(4)** two images with a one frame interval, and **(5) three images with one frame interval.**

Interpretability **We Apply Eigen-CAM (PCA method;** muhammad et al., 2020**) to positive cases** → heatmaps of where the model attends.

Data & Labeling

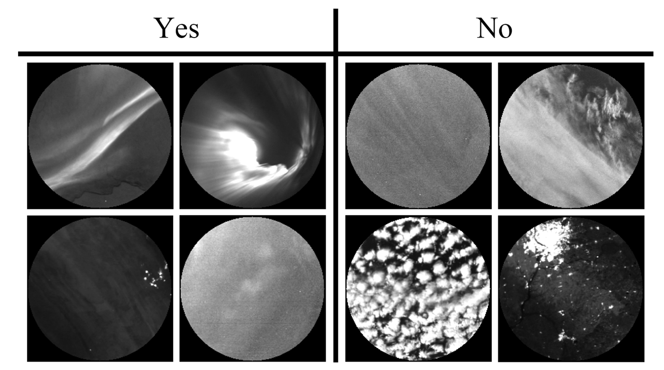

Figure 2. "Yes": discrete arcs, flaming, and diffuse auroras. "No": aurora-like clouds and city lights.

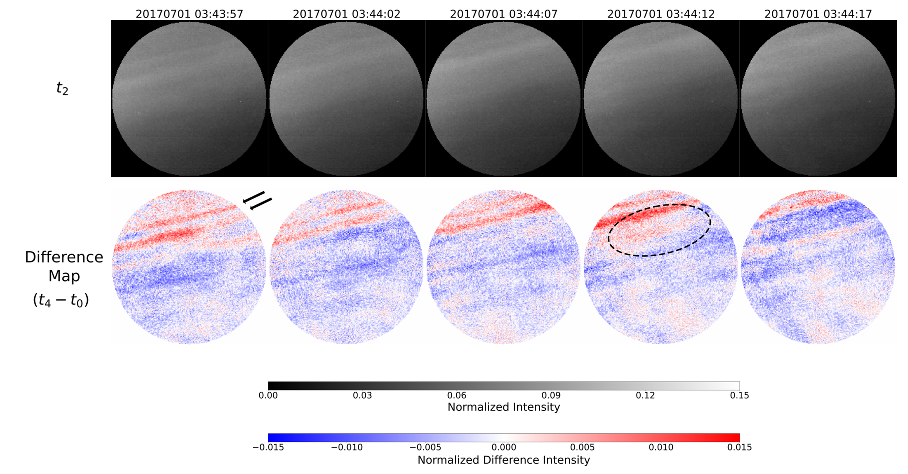

Figure 3. Visualization of input sequences and their corresponding difference maps between t_0 and t_4 for an auroral event predicted as "Yes".

Instrument **e-POP/FAI** (cogger et al., 2015) **NIR channel on CASSIOPE (polar orbit).**

Image type **Nadir-pointing** 256×256 frames, intensity in Rayleighs (R), 26° field of view, ~0.4–2.5 km/pixel at the ~110 km E-region; 100 ms exposure, 1 s cadence.

Data Period 2015–2017; test set: Jan. & Jul. (balances seasonal & geographic background).

Quality control We discard saturated frames (disk median > 1100 kR, e.g., moonlit reflections) and dark empty-ocean frames (≤ 40 kR). **These data are excluded from both training and evaluation.**

Labeling Manual, multi-annotator, and cross-checked; temporal behavior (flickering/pulsating) and altitude parallax aid classification.

Figure 3. The two solid black arrows indicate band-type features for reference. The black dashed oval highlights a flickering feature relative to the image taken four seconds earlier.

Dataset **Train:** 22,012 Yes / 22,901 No **Test:** 4,934 Yes / 4,486 No (Sequences)

Results

Table 1. Performance metrics of all five input configurations on the test set (Jan & Jul 2015–2017).

Model Number	Input frames	Accuracy	F1-score	Precision	Recall	MCC	AUC
1	1 frame (t_0)	0.81	0.80	**0.89**	0.73	0.64	0.89
2	2 frames (t_0, t_1)	0.80	0.80	0.84	0.76	0.61	0.87
3	3 frames (t_0, t_1, t_2)	0.81	0.81	0.88	0.74	0.64	0.90
4	2 frames with gap (t_0, t_2)	0.81	0.81	0.86	0.77	0.63	0.89
5	3 frames with gap (t_0, t_2, t_4)	**0.84**	**0.84**	0.86	**0.82**	**0.67**	**0.91**

- Overall, our models show show performance, and our best of our model is **#5 (three images with a one-frame gap)**.
- In some cases, the **Model 5 detects diffuse aurora**, whereas Model 1 does not.
- This is because diffuse aurora is hard to distinguish from clouds in a single frame.
- **Parallax effects then help,** since the difference in apparent velocity between the aurora and quasi-stationary terrestrial features (e.g., clouds) serves as a robust discriminant for the model.
- We use k-fold cross-validation (k = 5) to assess the robustness of our results and quantify uncertainty, and the results are robust across folds.
  (**Model 5** reaches 0.86 ± 0.02 (Acc) and 0.85 ± 0.02 (F1); std ≤ 0.03 confirms stability.)

Eigen-CAM Auroral Localization & Analysis

Eigen-CAM, a post-hoc, PCA-based interpretability technique, is applied to the outputs of Model 5.

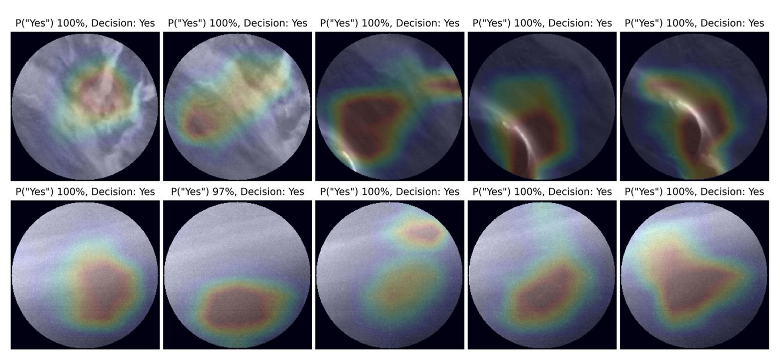

Figure 4. Eigen-CAM with true positive (**TP**) cases. The heatmaps overlaid on the images indicate regions of interest, where red areas correspond to regions the model focuses on most. Shown at the top of each image is the softmax probability of the "Yes" class, and binary decision (threshold at 0.5).

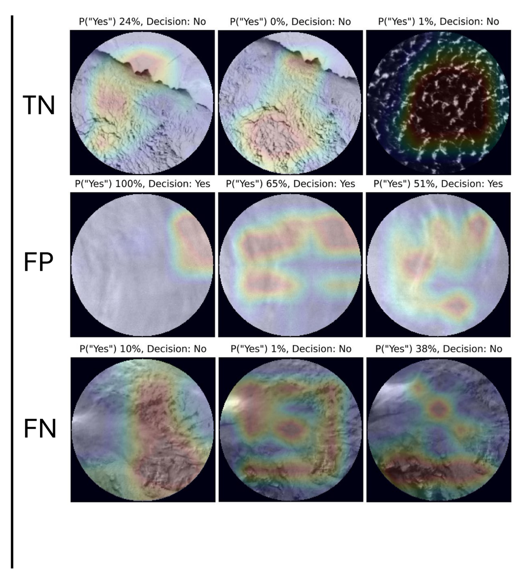

Figure 5. Eigen-CAM with true negative (**TN**, row 1), false positive (**FP**, row 2), false negative (**FN**, row 3) cases. The heatmaps overlaid on the images indicate regions of interest, where red areas correspond to regions the model focuses on most. Shown at the top of each image is the softmax probability of the "Yes" class, and binary decision (threshold at 0.5).

- Eigen-CAM is class-agnostic, so the resulting heatmaps highlight the most dominant (high-variance) features in the feature maps.
- In Figure 4. row 1, the Eigen-CAM method can extracts the newly-appearing auroral features.
- **Even in diffuse cases**, also following in Figure 3 cases, this indicating that the model attends to the flickering motions of the aurora.
- **Notably, although we do not label the location of the aurora, the model still highlights it in most True Positive (TP) cases.**
- In TN cases, the heatmaps focus on ground features (glaciers and clouds); in FP cases, they focus on band-like structures mistaken for aurora; and FN cases sometimes occur when the aurora covers only a small part of the frame.

Takeaways

- We present, to our knowledge, **the first deep-learning approach for classifying auroras in near-infrared (NIR) satellite** time-series images from e-POP/FAI.
- It is **difficult to distinguish auroral emissions from other bright features**, such as clouds and city lights, **however, our model can classify these frames.**
- The best model is **#5 (three images with one-frame gap)**, reaching high accuracy = 0.84, F1 = 0.84 and AUC = 0.91.
- **Temporal context is the key driver for the model,** it captures pulsating dynamics (~2–20 s) and altitude parallax effect.
- Note that, Eigen-CAM is class-agnostic, so it cannot fully segment the aurora and occasionally highlights non-auroral regions.
- **However, the model mostly highlights the auroral features in True Positive (TP) cases, even we do not label the location of the aurora.**
- In operation point, **combining Eigen-CAM with the geolocation** stored in each frame's metadata lets the "Yes" detections **provide candidate auroral-oval coordinates along the orbit.**
- The framework transfers directly to other missions such as the Republic Of Korea Imaging Test System (ROKITS), supporting nowcasting of the auroral-oval boundary, and future work will add a dedicated segmentation model.

Junmu Youn

Our code is available on GitHub. →

Github / CV

CEDAR Workshop 2026 Des Moines, Iowa · Dataset: doi.org/10.5281/zenodo.16974375 · Code: github.com/JunmuYOUN/SpaceAI_epop_Aurora · Data archive: epop-data.phys.ucalgary.ca