Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

SOLARIS: Solar Activity Recognition and Interpretation System

Mario A. Jara
Universidad Autónoma de Nuevo León
Axel F. Rodriguez Vazquez
Universidad Autónoma de Nuevo León
Carlos A. Guerrero Peña
carlos.guerrerop@uanl.edu.mx

Universidad Autónoma de Nuevo León
Esmeralda Romero-Hernandez
Universidad Autónoma de Nuevo León
Juan José Tapia Armenta
Centro de Investigación y Desarrollo de Tecnología Digital (CITEDI-IPN)
Daysi Martínez
Universidad Autónoma de Nuevo León
Yizzel Ibarra
Universidad Autónoma de Nuevo León
Sandra Ayala
Universidad Autónoma de Nuevo León
Mario Rodríguez-Martínez
Universidad Nacional Autónoma de México
David Estrada
Universidad Autónoma de Nuevo León
G. Pérez-Tijerina
Universidad Autónoma de Nuevo León

Research Article

Keywords: solar cycle, sunspots, active regions, machine-learning, detection

Posted Date: May 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9596586/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

SOLARIS: Solar Activity Recognition and Interpretation System

Mario A. Jara[1] · Axel F. Rodriguez Vazquez[1] · Carlos A. Guerrero Peña[1] · Esmeralda Romero-Hernandez[1] · Juan José Tapia Armenta[2] · Daysi Martínez[1] · Yizzel Ibarra[1] · Sandra Ayala[1] · Mario Rodríguez-Martínez[3] · David Estrada[1] · G. Pérez-Tijerina[1]

© **The author(s)** ●●●●

✉ C. A. Guerrero Peña
carlos.guerrerop@uanl.edu.mx

M. A. Jara
mario.jarare@uanl.edu.mx

A. F. Rodriguez Vazquez
axel.rodriguezvzqz@uanl.edu.mx

E. Romero-Hernandez
esmeralda.romerohdz@uanl.edu.mx

J. J. Tapia Armenta
jjtapia@citedi.mx

D. Martínez
1808011daysimtz@gmail.com

Y. Ibarra
pa.yizzel.02@gmail.com

S. Ayala
sandra.ayalagmz@uanl.edu.mx

M. Rodríguez-Martínez
mrodriguez@enesmorelia.unam.mx

D. Estrada
david.estradasnt@uanl.edu.mx

Eduardo Pérez-Tijerina
eduardo.pereztj@gmail.com

[1] Universidad Autónoma de Nuevo León, Facultad de Ciencias Físico Matemáticas, Monterrey, NL, México

[2] Centro de Investigación y Desarrollo de Tecnología Digital, Tijuana, BC, México

[3] Escuela Nacional de Estudios Superiores, Unidad Morelia, Universidad Nacional Autónoma de México, Morelia, Michoacán, México

Abstract
The automated detection and classification of solar activity features—such as sunspots, coronal holes, and prominences—is essential for understanding the magnetic dynamics of the Sun and for supporting space-weather monitoring. In this work, we introduce SOLARIS: Solar Activity Recognition and Interpretation System, a deep-learning system for real-time identification of solar features based on the YOLO26 object-detection framework. The model was trained on more than 12,000 full-disk SDO/HMI continuum intensitygrams and optimized through Ray Tune hyperparameter search. On a held-out test set, SOLARIS achieves a mean Average Precision at IoU=0.5 (mAP@0.5) of 0.930, with per-class scores of 0.967 for sunspots, 0.959 for coronal holes, and 0.864 for prominences, demonstrating strong generalization across diverse solar morphologies. Beyond its research applications, SOLARIS is being deployed operationally at the LANCE–UANL Solar Observatory, where it serves as the core detection engine for a new real-time monitoring program based on ground-based imaging. This implementation provides an operational testbed for validating the model under ground-based conditions, enabling continuous detection of solar atmospheric features and supporting future developments in tracking, education, and forecasting. By integrating practical deployment, SOLARIS advances both heliophysics research and the development of accessible, low-latency solar-monitoring infrastructure.

Keywords: solar cycle, sunspots, active regions, machine-learning, detection

1. Introduction

The solar atmosphere is a highly dynamic environment, constantly shaped by magnetic activity that unfolds across a broad range of spatial and temporal scales. Phenomena such as sunspots, active regions, prominences, and coronal holes are all surface-level expressions of deeper magnetic processes driven by the Sun's 22-year magnetic cycle (Hathaway, 2015; Wiegelmann, Thalmann, and Solanki, 2014). Understanding and detecting these features is not only a fundamental challenge in heliophysics but also a crucial task for space-weather monitoring, since solar activity can affect satellites, communication and navigation systems, aviation, and even ground-based technologies such as power grids (Pulkkinen, 2007; Schrijver, 2015; Opgenoorth et al., 2019; Kauristie et al., 2021). In this context, long-term monitoring of solar activity also contributes to the characterization and prediction of solar-cycle variability (Petrovay, 2020).

Sunspots are among the most visible and historically studied features of solar magnetic activity. Their morphology and magnetic complexity are closely tied to transient energetic events such as solar flares and coronal mass ejections (CMEs) (Toriumi and Wang, 2019; Schrijver, 2007), making them key indicators of evolving solar conditions. The spatial distribution, size, and topology of sunspots carry predictive potential regarding the emergence and evolution of solar active regions. Coronal holes and prominences, meanwhile, trace open and closed magnetic configurations of the corona, respectively, linking photospheric dynamics to solar wind structures and eruptive behavior.

The continuous observation of the solar disk has been revolutionized by the *Solar Dynamics Observatory* (SDO; Pesnell, Thompson, and Chamberlin 2012), which provides uninterrupted, high-resolution full-disk imaging across multiple wavelengths. The Helioseismic and Magnetic Imager (HMI) instrument onboard SDO delivers intensitygrams and magnetograms suitable for studying photospheric dynamics with a cadence of one image per minute. However, the immense and continuously growing volume of solar observations, which can reach terabytes per day, makes the manual identification and classification of activity features increasingly impractical. This motivates the development of automated pipelines capable of detecting and classifying solar features in real time, with minimal human intervention.

Complementary to SDO's high-cadence observations, Sun's long-term datasets have recently expanded the temporal baseline of solar variability studies. For instance, Sandford et al. (2025) presented 28 years of extreme ultraviolet (EUV) light curves from the *Solar and Heliospheric Observatory*/Extreme ultraviolet Imaging Telescope (SOHO/EIT) instrument, providing valuable insights into the evolution of irradiance, rotation modulation, and activity-cycle dependencies. This extensive dataset provides the empirical framework for linking global irradiance trends with high-resolution morphological analyses. Achieving this connection requires an automated feature-recognition pipeline, such as SOLARIS (*Solar Activity Recognition and Interpretation System*).

Recent developments in this field highlight the rapid evolution of deep-learning approaches for solar monitoring. For example, Gong et al. (2024) introduced *ReDetGraphTracker*, a model designed to detect and track active regions in SDO/HMI magnetograms with enhanced temporal coherence and mean average precision. These efforts have demonstrated the potential of object-detection networks to reliably identify and follow magnetic features across multiple observing campaigns.

Another solid contribution in this direction was achieved by Pan et al. (2025), who developed the *Solar Active Region Detection* (SARD) system, a YOLOv8-based architecture optimized for identifying active regions (ARs) in line-of-sight magnetograms from SDO/HMI covering Solar Cycle 24. While SARD represents a significant advancement toward the automation of magnetic feature recognition, its scope remains limited to a single data modality and a single class of features. In contrast, the present work extends this framework by introducing SOLARIS, a YOLO26-based model trained on continuum intensitygrams from SDO/HMI to simultaneously detect multiple solar features, including sunspots, coronal holes, and prominences, under the more complex activity conditions of Solar Cycle 25.

SOLARIS leverages an attention-centric backbone with residual feature aggregation and employs Ray Tune hyperparameter optimization to achieve a balanced trade-off between accuracy and computational efficiency. Although predictive forecasting is beyond the present scope, this study focuses on establishing the technical foundation, detection accuracy, and practical applicability of the proposed framework for automated solar feature recognition.

The remainder of this paper is organized as follows. Section 2 describes the methodology, including the dataset, preprocessing steps, labeling strategy, and

the YOLO26 model architecture. Section 3 presents the training procedure and the hyperparameter optimization performed with Ray Tune. Section 4 reports the model performance and evaluation results. Finally, Section 5 summarizes the main conclusions and outlines future developments of the SOLARIS system.

2. Methodology

2.1. Dataset and Observations

The dataset used in this study consists of 12,385 full-disk continuum intensitygrams obtained by the *Helioseismic and Magnetic Imager* (HMI) onboard the SDO, between May and August 2024. This period corresponds to the ascending phase of Solar Cycle 25, characterized by an increment of magnetic activity and the recurring manifestation of complex sunspot groups, prominences, and coronal holes. The images were retrieved at a spatial resolution of 4096×4096 pixels and a temporal cadence of 45 s, then resampled to 1024×1024 to reduce computational load during model training.

All images were pre-filtered to remove frames with telemetry loss or significant cosmic-ray artifacts using quality flags provided by the Joint Science Operations Center (JSOC). The selected sample ensures a homogeneous distribution in heliographic longitude and time, providing a representative set of diverse solar morphologies. The truth labels for sunspots, coronal holes, and prominences were manually generated using a combination of NOAA/SWPC solar region maps and visual verification on SDO/AIA composite images (171, 193, and 304 Å).

2.2. Preprocessing and Limb-Darkening Correction

To minimize the photometric gradient across the solar disk, we implemented a limb-darkening correction step using the *Solar Limb Darkening Toolkit* (SLDTk; Ladegaard, Wingham, and Heron 2017). Each image was first normalized by its mean disk intensity, then corrected using a fifth-order polynomial fit to the radial intensity profile:

$$I(\mu) = I_0 \left(1 - \sum_{k=1}^{5} a_k (1 - \mu^k) \right), \tag{1}$$

where $\mu = \cos\theta$ and θ is the heliocentric angle. Typical limb-darkening coefficients a_k were derived followingg Moon et al. (2017), yielding a consistent center-to-limb intensity variation across the entire dataset.

After correction (see Figure 1), images were contrast-enhanced using adaptive histogram equalization (CLAHE) to improve feature visibility, and background levels were standardized via median normalization. Additional steps included (i) masking of the off-disk region using an automatic circular Hough transform, and (ii) data augmentation through random rotations ($\pm 20°$), scaling (90–110%), and intensity jittering ($\pm 5\%$), ensuring rotational invariance of the detection model.

All continuum images used in this study were obtained from the JSOC using the `hmi.ic_45s` series, which provides continuum intensitygrams at a cadence

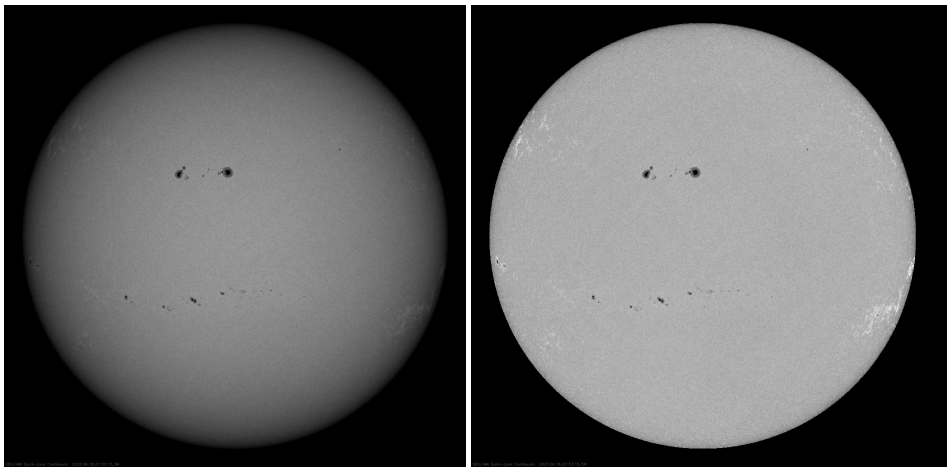


Figure 1. a) Full-disk photosphere image. b) Same image with the limb darkening correction

of 45 s. The dataset spans the interval from 1 May 2024 to 31 August 2024, corresponding to the ascending phase of Solar Cycle 25. Data were retrieved through the JSOC Export API[1].

To homogenize temporal sampling, one intensitygram was selected every 6 hours, producing 12,385 full-disk images used for training, validation, and testing. Each FITS file includes standard HMI metadata (WCS headers, pointing, and calibration keywords), ensuring full reproducibility of the dataset.

2.3. Data Structure and Labeling Strategy

Each preprocessed image was stored alongside its corresponding annotation file in YOLO format, containing bounding-box coordinates and class identifiers (0: sunspot, 1: coronal hole, 2: prominence). The dataset was divided into 70% training, 15% validation, and 15% test subsets, ensuring temporal separation between sets to avoid overfitting on near-identical solar configurations. This organization follows a modular data structure tested in an earlier prototype system that we developed, facilitating scalability for future class extensions (e.g., filaments, plages, etc.).

The labeling of sunspots followed the morphological and magnetic criteria established in the NOAA catalogues, where an active region is characterized by the presence of one or more umbrae surrounded by penumbral structure and classified according to its magnetic configuration (simple α, bipolar β, or complex $\beta\gamma\delta$ types). For each HMI continuum frame, bounding boxes were drawn to encompass both the umbra and penumbra of each identified group, ensuring that nearby pores were included only when magnetically connected to the main region, as verified through the corresponding HMI line-of-sight magnetograms.

Coronal holes were annotated using the intensity contrast method applied to SDO/AIA 193 Å images, which are highly sensitive to emission from Fe XII in

[1] http://jsoc.stanford.edu/

the corona at ∼1.3 MK. Candidate regions were extracted by adaptive thresholding and verified by their unipolar magnetic signature in co-aligned HMI magnetograms, following the methodology of Lowder, Qiu, and Leamon (2017). Prominence contours were derived from AIA 304 Å data, where these structures appear as bright, elongated filaments against the darker chromospheric background (Parenti, 2014). Only stable prominences with lifetimes exceeding one hour were retained, excluding transient brightenings and short-lived surges.

All labels were manually reviewed to ensure consistency. Discrepancies greater than 5% in bounding-box area or position were re-evaluated jointly, and a final mask was produced. Special attention was given to projection effects near the solar limb, where apparent feature sizes can be distorted; in such cases, manual correction was applied based on helioprojective geometry using World Coordinate System (WCS) metadata (Thompson, 2006; Müller et al., 2017). This validation process yielded a final annotated set of 12,385 images with high spatial and temporal consistency across all three feature classes. They were divided into 8,670 images (70 %) for training, 1,858 (15 %) for testing, and 1,858 (15 %) for validation.

2.4. Model Architecture: YOLO26 Framework

The detection model implemented within the SOLARIS framework is based on the YOLO26 architecture (Jocher and Qiu, 2026), a real-time object detector that integrates attention-centric feature extraction with enhanced multi-scale fusion. The backbone employs the Residual Efficient Layer Aggregation Network (R-ELAN), whose purpose is to deepen feature aggregation while maintaining computational efficiency. Given an input feature map X, each convolutional transformation in R-ELAN is described as follows:

$$\mathrm{F_i} = \sigma(\mathrm{W_i} * \mathrm{X} + \mathrm{b_i}), \tag{2}$$

where W_i and b_i denote the convolutional weights and biases, and $\sigma(\cdot)$ is the nonlinear activation. This residual formulation improves gradient flow and stabilizes training across the multiple spatial scales required for solar-feature detection.

A local attention mechanism, implemented through Area Attention, enhances the model's sensitivity to compact and high-contrast structures such as small sunspots and limb prominences. This mechanism computes attention weights through:

$$\mathrm{A_{ij}} = \frac{\exp(\mathrm{Q_i} \cdot \mathrm{K_j})}{\sum_{\mathrm{k}} \exp(\mathrm{Q_i} \cdot \mathrm{K_k})}, \tag{3}$$

where Q and K represent the queries and keys derived from the incoming feature maps. By combining R-ELAN with local attention, YOLO26 improves discriminative power in regions of the solar disk affected by steep intensity gradients or foreshortening.

Bounding-box prediction follows the standard YOLO regression parameterization:

$$\hat{\mathrm{b}} = (\hat{\mathrm{x}}, \hat{\mathrm{y}}, \hat{\mathrm{w}}, \hat{\mathrm{h}}, \hat{\mathrm{c}}), \tag{4}$$

where $(\hat{x}, \hat{y})$ is the predicted box center, $(\hat{w}, \hat{h})$ its width and height, and $\hat{c}$ the confidence score. During training, the loss function combines localization, classification, and confidence terms:

$$\mathcal{L} = \lambda_{\text{loc}}\, \mathcal{L}_{\text{IoU}} + \lambda_{\text{cls}}\, \mathcal{L}_{\text{cls}} + \lambda_{\text{conf}}\, \mathcal{L}_{\text{conf}}, \tag{5}$$

where $\mathcal{L}_{IoU}$ denotes a generalized IoU-based penalty (Rezatofighi et al., 2019), $\mathcal{L}_{cls}$ is the cross-entropy or focal-loss classification term, and $\mathcal{L}_{conf}$ regulates the predicted confidence of each box.

To tailor YOLO26 to our solar dataset, the training hyperparameters were optimized using a dedicated search strategy (see Section 2.5). This process aims to identify a configuration that maximizes detection accuracy while preserving real-time inference capability.

YOLO26 has demonstrated competitive performance in real-time object detection tasks, achieving high accuracy and low latency on modern GPU architectures (Jocher and Qiu, 2026). These characteristics make it particularly suitable for detecting solar features with diverse spatial scales and morphologies, as required in the SOLARIS framework.

2.5. Training procedure and hyperparameter optimization

Hyperparameter optimization was conducted using the Ray Tune framework (Liaw et al., 2018), a scalable library designed for distributed hyperparameter search. In this stage, a fixed YOLO26 base architecture was used across all experiments, while only the training hyperparameters were varied. This ensures that performance differences arise exclusively from the optimization of the training configuration rather than architectural changes.

Each sampled configuration corresponds to an independent training run of the same base model. For every trial, the model was trained using the PyTorch implementation of YOLO26 on a workstation equipped with three NVIDIA TITAN RTX GPUs (24 GB VRAM each) and 512 GB RAM. Training was performed for up to 50 epochs with a batch size of 32 and an input resolution of 640×640 pixels. Optimization was carried out using stochastic gradient descent (SGD) with automatic mixed precision (AMP) enabled to improve computational efficiency. An early-stopping criterion with a patience of 10 epochs was applied to prevent unnecessary computation.

Model performance was evaluated using the mean average precision across multiple IoU thresholds (mAP50–95), which served as the objective metric for hyperparameter selection.

To efficiently explore the search space, Ray Tune was combined with the Asynchronous Successive Halving Algorithm (ASHA) (Li et al., 2018), which dynamically allocates computational resources by promoting promising trials and terminating underperforming ones at early stages. This strategy significantly reduces the overall computational cost while maintaining an effective exploration of the parameter space.

The search space included key training hyperparameters such as the initial learning rate, learning-rate factor, momentum, weight decay, warmup epochs,

and loss-balance parameters. A total of 40 trials were executed, each corresponding to a different configuration sampled from the defined ranges.

The best-performing trial achieved a validation score of mAP50–95 = 0.5629, indicating a balanced performance between localization accuracy and classification across multiple scales. The corresponding optimal hyperparameters are summarized in Table 1.

Table 1. Hyperparameter search space and optimal configuration obtained using Ray Tune.

Hyperparameter	**Search range**	**Optimal value**
Learning Rate	$[1 \times 10^{-5},\ 5 \times 10^{-2}]$	0.023
Learning Rate Factor	$[0.001,\ 0.050]$	0.046
Momentum	$[0.750,\ 0.950]$	0.837
Weight Decay	$[1 \times 10^{-6},\ 5 \times 10^{-4}]$	0.000
Warmup Epochs	$[0.000,\ 4.000]$	1.393
Box Loss Gain	$[0.030,\ 0.200]$	0.161
Classification Loss Gain	$[0.300,\ 2.500]$	0.338

In addition to the best-performing configuration, the distribution of the top-ranked trials reveals a consistent convergence behavior. As shown in Figure 2, the top 10 configurations exhibit a relatively narrow performance range, with only moderate variations in mAP50–95. This suggests that the optimization process effectively identified a region of high-performing solutions within the hyperparameter space, rather than relying on a single isolated optimum. Such consistency supports the robustness of the training procedure and confirms the effectiveness of the exploration strategy.

The configuration identified in this stage was selected as the final training setup. The complete training procedure using these optimized hyperparameters is described in the following section.

2.6. Training performance and evaluation

The performance of SOLARIS during training was monitored through a range of diagnostic metrics. The first set of indicators is shown in Figure 3, which presents the F1 curve and the Precision–Recall (PR) curves for each class. The F1 curve (Figure 3a) illustrates the balance between precision and recall as a function of the confidence threshold. The model reaches its maximum F1 score of 0.89 at a threshold of 0.454, indicating an optimal operating point that minimizes both false positives and false negatives.

The PR curves displayed in Figs. 3b–c provide class-specific insight. During validation, the model achieves a mean mAP@0.5 of 0.935, with per-class values of 0.974 for *sunspot*, 0.962 for *coronal_hole*, and 0.868 for *prominence*. The corresponding precision–confidence behavior shows that precision approaches 1.00 at high confidence thresholds ($c \approx 0.963$), indicating that the model produces highly reliable predictions when stricter confidence filtering is applied.

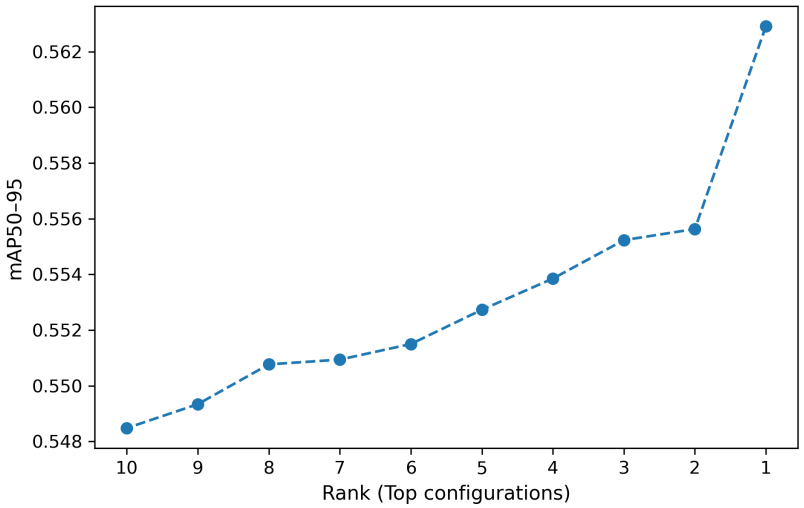


Figure 2. Performance of the top 10 hyperparameter configurations ranked by mAP50–95. The observed variation across configurations suggests that the optimization process explored a region of high-performing solutions within the hyperparameter space.

To further characterize performance, Figure 4 shows the confusion matrices for validation (normalized) and test (absolute counts). The model demonstrates strong discriminability for *sunspot* and *coronal_hole*, which achieve the highest classification accuracy. The *prominence* class exhibits comparatively higher confusion, primarily due to morphological similarities with surrounding structures and background regions.

The overall evolution of losses and evaluation metrics is summarized in Figure 5. Both training and validation losses, including `box_loss`, `cls_loss`, and `dfl_loss`, decrease smoothly throughout the training process without signs of divergence. At the same time, performance metrics such as precision, recall, mAP@0.5, and mAP50–95 show consistent improvement and stabilization. In particular, recall reaches values close to 0.878, while mAP50–95 converges to approximately 0.573. These results indicate a stable training process and confirm that SOLARIS effectively captures multi-scale solar features across different morphological regimes.

3. Results

3.1. Postprocessing and Output Generation

After inference, raw detections are refined using non-maximum suppression (NMS) with an IoU threshold of 0.45 and a confidence filter of 0.3. This procedure removes redundant bounding boxes and ensures that each physical structure is represented by a single detection. For visualization, SOLARIS overlays the

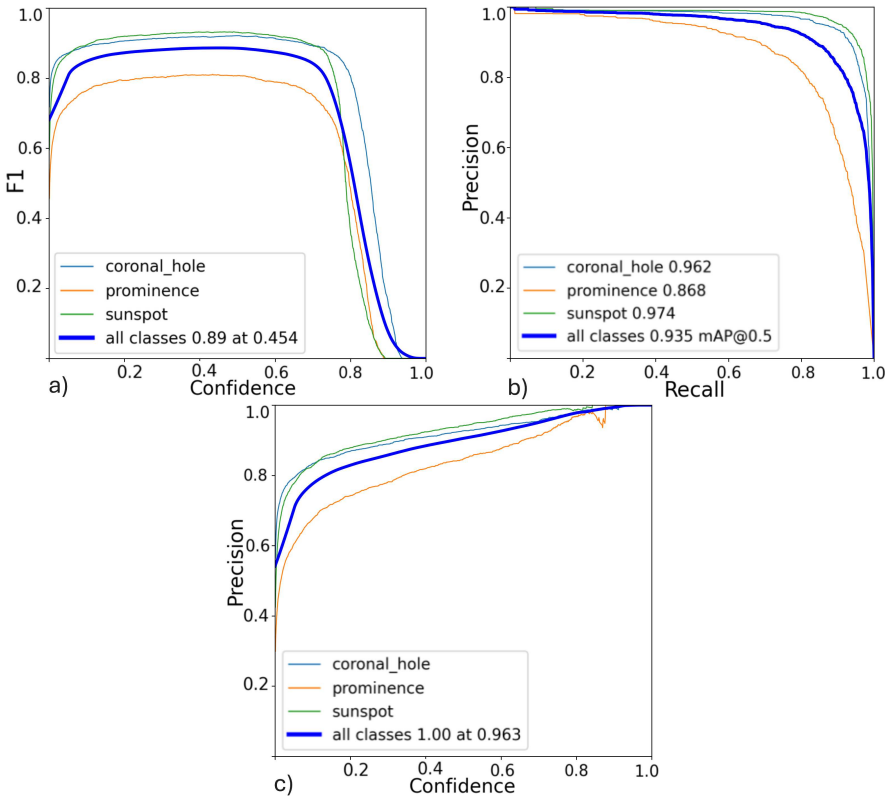


Figure 3. (a) F1 score as a function of the confidence threshold, with a maximum of 0.89 at $c = 0.454$. (b) Precision–Recall curves per class (mAP@0.5 = 0.935). (c) Precision as a function of confidence, reaching values close to 1.00 at high thresholds.

predicted bounding boxes and class labels onto the original full-resolution HMI continuum frames.

To obtain a sharper morphological description of each detected feature, an optional segmentation step is applied. A watershed-based algorithm is executed inside the predicted bounding box, following the morphometric procedure validated during model development. Although segmentation is not required for object detection, these masks improve the extraction of structural descriptors—such as projected area, elongation, or radial extension—which are relevant for future statistical studies of solar activity.

Each processed frame produces two outputs: (1) a graphical overlay showing detections, and (2) a metadata file containing bounding-box coordinates, class ID, detection confidence, and the corresponding HMI timestamp. These products facilitate time-series analyses of feature evolution (e.g., sunspot growth or prominence dynamics) and can be integrated with multi-instrument data streams such as GOES flare catalogs or SDO/AIA sequences.

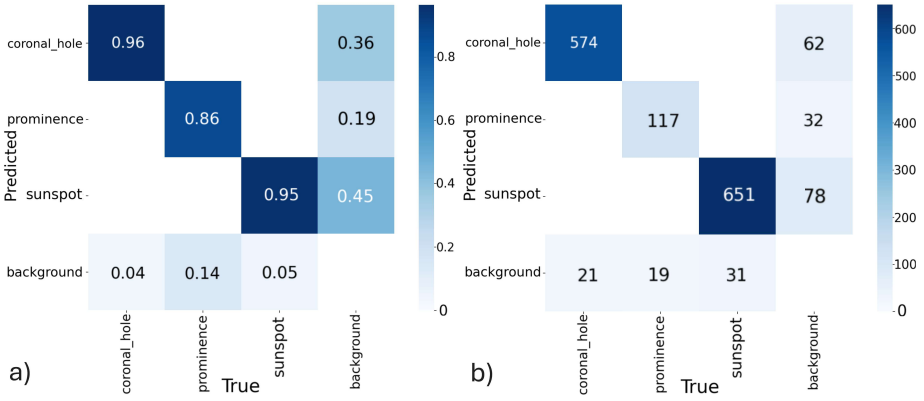


Figure 4. (a) Normalized confusion matrix for the validation set. (b) Absolute-count confusion matrix for the test set. *Sunspot* and *coronal_hole* exhibit strong discriminability, while *prominence* shows higher confusion.

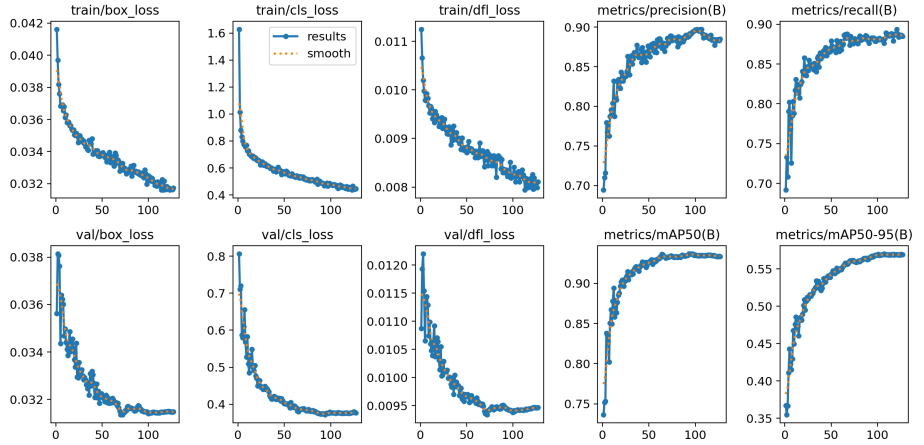


Figure 5. Training evolution across epochs. The figure shows the progression of training losses (`box_loss`, `cls_loss`, `dfl_loss`), validation losses, and evaluation metrics including precision, recall, mAP@0.5, and mAP50–95 as a function of the epoch index. The model converges stably, reaching maximum values of recall ≈ 0.878, precision ≈ 0.886, mAP@0.5 ≈ 0.930, and mAP50–95 ≈ 0.573.

3.2. Quantitative Performance on the Test Set

The model was evaluated on a held-out test set comprising 1857 images and 4551 labeled instances. SOLARIS achieves a mean precision of 0.886, recall of 0.878, and mAP@0.5 of 0.930 across all classes. The corresponding mAP@0.5–0.95 reaches 0.573, indicating robust performance across multiple IoU thresholds. Performance varies moderately among the three feature types:

- Sunspots: show the highest detectability ($P = 0.913$, $R = 0.935$, mAP@0.5 = 0.967), reflecting their strong contrast and compact morphology in continuum intensitygrams.

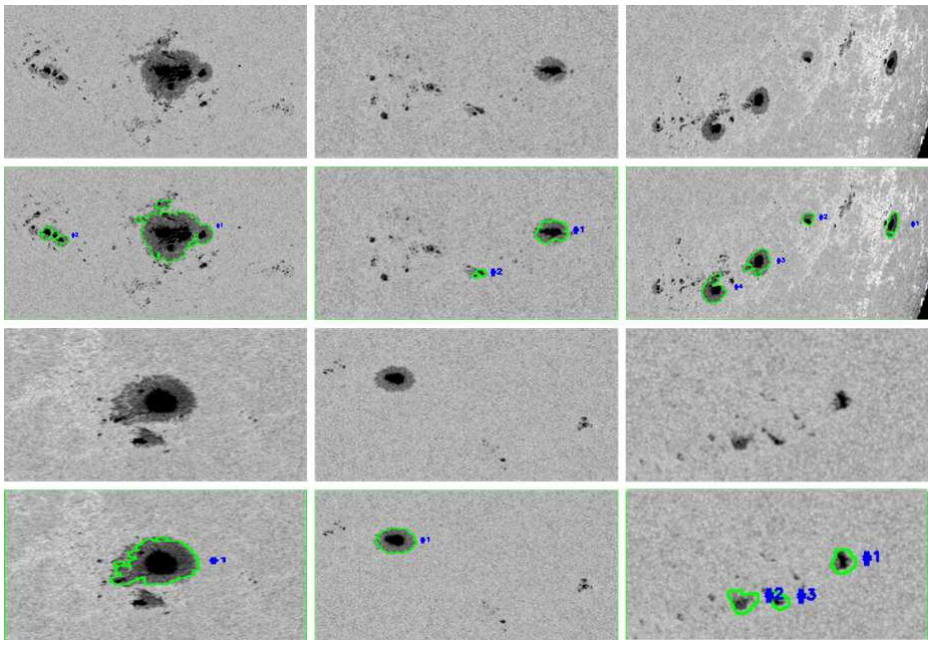


Figure 6. Examples of SOLARIS detections on HMI continuum images, focused on sunspot-bearing regions. The green contours show the watershed-based refinement applied within the predicted bounding boxes, enabling the extraction of umbral and penumbral morphology..

- Coronal holes: achieve similarly high performance ($P = 0.917$, $R = 0.903$, mAP@0.5 $= 0.959$), consistent with their coherent spatial extent despite their lower contrast.
- Prominences: being more filamentary and located near the limb, remain the most challenging class ($P = 0.829$, $R = 0.797$, mAP@0.5 $= 0.864$).

These differences highlight the influence of morphological complexity and limb-projection effects on detection accuracy. The lower mAP@0.5–0.95 for prominences further reflects their sensitivity to localization errors under stricter IoU thresholds. Table 2 summarizes the test set metrics. The model achieves high precision, recall, and mAP, confirming its robustness. As shown in Figure 6, SOLARIS successfully identifies multiple features across the solar disk.

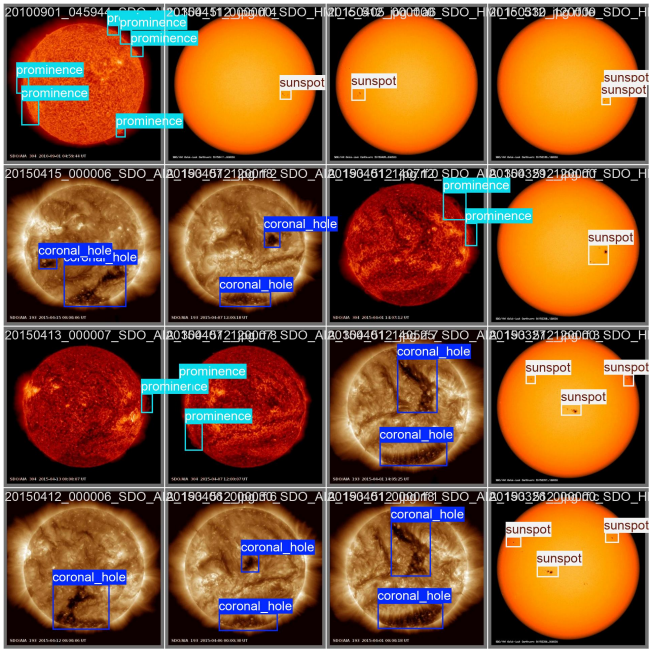
Table 2. Evaluation metrics on the test set.

Class	Images	Instances	Box(P)	R	mAP50	mAP50-95
all	1857	4551	0.886	0.878	0.930	0.573
coronal_hole	790	1535	0.917	0.903	0.959	0.652
prominence	420	1322	0.829	0.797	0.864	0.489
sunspot	647	1694	0.913	0.935	0.967	0.578

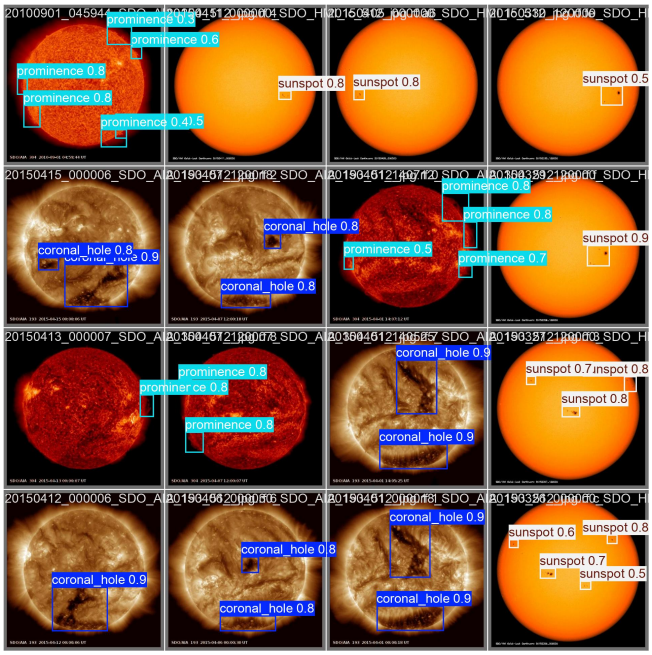
3.3. Qualitative Evaluation

Figure 7 compares the manually annotated ground truth with the predictions generated by SOLARIS. The model successfully identifies most instances even in images with complex active-region topologies or multiple overlapping features. Coronal holes are robustly detected both near disk center and toward the limb, while sunspots are identified accurately regardless of size. Prominence detections tend to cluster near the limb, where projection effects are strongest, but remain consistent with the expected shapes and positions.

The qualitative results confirm the reliability of the bounding-box predictions and suggest that SOLARIS can be used for further downstream applications, including catalog construction, tracking, or statistical activity profiling.



a)



b)

Figure 7. (a) Ground truth labels with manual annotations. (b) YOLO26 model predictions including detected classes, bounding boxes, and confidence levels. Panels include the auxiliary AIA views used during label verification; model inference, however, was performed on the corresponding HMI continuum frames.

3.4. Confusion Analysis and Class Behavior

The confusion matrices displayed in Figure 4 illustrate that most misclassifications involve confusion between prominences and coronal-hole boundaries. This behavior is physically reasonable: limb-adjacent regions often exhibit intensity gradients or partial occultation effects that resemble diffuse prominence edges.

However, despite these challenges, over 88% of all instances across the test set are correctly classified. The consistency between validation and test performance suggests that the feature representations learned by the model generalize effectively across different solar states sampled between May and August 2024.

3.5. Training Evolution and Convergence

The training evolution shown in Figure 5 demonstrates stable learning dynamics. Training and validation losses decrease smoothly with no evidence of divergence or oscillations, indicating that the hyperparameter configuration obtained via Ray Tune search enabled efficient and robust convergence.

Recall and precision rise consistently, reaching ≈ 0.88 and ≈ 0.89 respectively, while mAP@0.5 stabilizes near 0.93 and mAP50–95 near 0.57. The latter value is characteristic of multi-scale solar scenes containing small, morphologically diverse objects, where high-IoU predictions are more difficult to achieve. Overall, these results confirm that YOLO26, combined with limb-darkening correction and Ray Tune-based optimization, provides a competitive detection framework for solar feature analysis.

4. Discussion

Our results demonstrate that real-time detection of multiple solar features is feasible using a unified deep-learning model trained solely on HMI continuum intensitygrams. The performance achieved by SOLARIS compares favorably with existing single-feature detectors, while extending the capability to multiple classes simultaneously. This multi-class approach is essential for constructing holistic descriptions of solar activity and for supporting downstream applications such as flare precursors or active-region evolution studies.

To place the performance and scope of SOLARIS in context, Table 3 summarizes representative recent deep-learning approaches for solar-feature detection and tracking. As can be seen, SOLARIS differs from previous approaches by addressing a multi-class detection problem using continuum intensitygrams rather than magnetograms alone. Although this choice imposes limitations for structures whose primary signatures are thermal or magnetic, such as coronal holes and prominences, the results demonstrate that a unified detection framework can recover useful multi-feature information from a single observational modality.

A notable strength of the system is its robustness across a wide variety of solar conditions sampled during the rising phase of Solar Cycle 25. The model performs well even on images containing complex sunspot groups, scattered coronal-hole fragments, or faint limb prominences, illustrating its adaptability to the diverse morphologies present in the solar photosphere and chromosphere.

Table 3. Comparison between SOLARIS and representative recent deep-learning approaches for solar-feature detection and tracking.

Study	Input data and target	Task and metric	Main limitation
Gong et al. (2024)	SDO/HMI magnetograms; active regions	Detection/tracking; high accuracy and temporal coherence	Focused on magnetic active regions; not intended for multi-class atmospheric-feature detection.
Pan et al. (2025)	SDO/HMI LOS magnetograms; active regions	Active-region detection; AP $\approx 94\%$	Restricted to one feature class and one magnetic-data modality.
This work	SDO/HMI continuum intensitygrams; sunspots, coronal holes, prominences	Multi-class detection; mAP@0.5 = 0.930, mAP@0.5–0.95 = 0.573	Limited direct continuum contrast for coronal holes and prominences, which are better traced in EUV or magnetic data.

Because coronal holes and prominences are not primarily continuum features, their inclusion in the HMI-continuum training set should be interpreted as cross-modality labeling: the model learns continuum-contextual signatures associated with structures identified in AIA EUV channels, rather than direct EUV emission morphology.

4.1. Operational Deployment and Real-Time Monitoring

In parallel with the development of the present model, our research group is implementing the SOLARIS pipeline as the foundation of a real-time monitoring program for photospheric solar activity at the LANCE–UANL Solar Observatory. The YOLO26-based detector will operate as the core identification module for sunspots, coronal holes, and prominences, enabling rapid alerts on evolving solar conditions.

As a preliminary external test under ground-based observing conditions, SOLARIS was applied to 250 Solar Hα (6562.8 Å) images obtained at the Solar Observatory of the Laboratorio de Ciencias GeoEspaciales (LACIGE, UNAM; Lat. N 19.65, Long. W 101.22). These observations were acquired with a SolarMax 90 mm Meade refractive telescope equipped with a ZWO ASI183MM monochromatic camera, producing images with a resolution of 5496×3672 pixels. The LACIGE images were not part of the SDO/HMI training, validation, or test sets; therefore, they provide an initial domain-shift assessment of the pipeline using independent ground-based data. Because the analyzed Hα images were centered on sunspot-bearing regions, this test was restricted to sunspot detections.

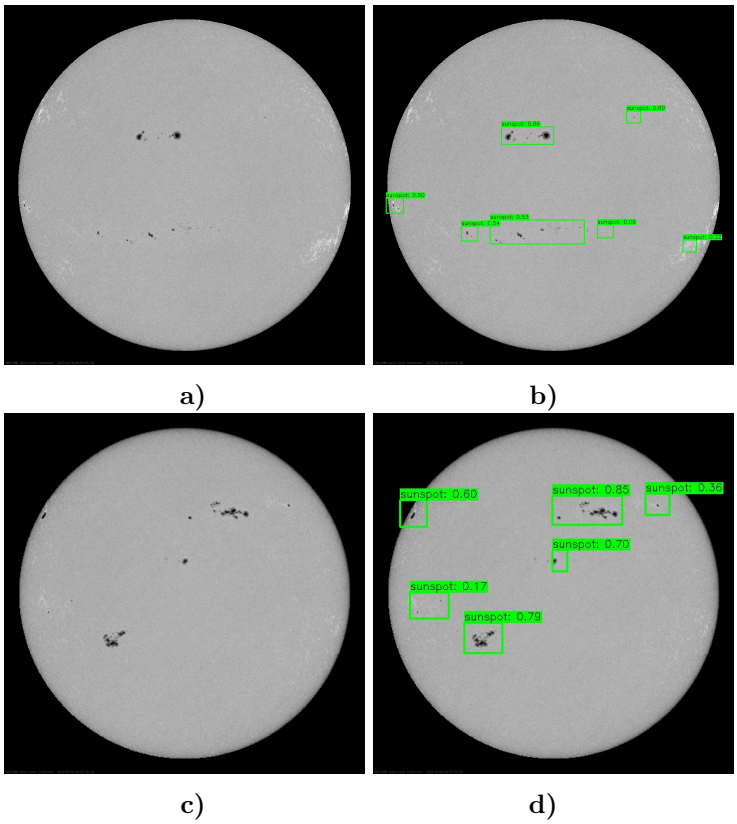


Figure 8. (a) & (c) Solar images showing sunspots, obtained at the ground-based Solar Observatory of the LACIGE, UNAM. (b) & (d) Same images after the SOLARIS identification process.

Table 4. Detection-level statistics from the external LACIGE Hα image test. The bounding-box area is given in image pixels.

Feature	Frames	Detections	Det./Frame	Mean Conf.	Std. Conf.	Mean Area
Sunspot	250	1348	5.392	0.559	0.257	3.823×10^3

Figures 8a and 8c show representative original images, while Figures 8b and 8d present the corresponding SOLARIS outputs. In addition to these qualitative examples, the pipeline stores each detection in a CSV metadata file containing the bounding-box coordinates, class ID, confidence score, and timestamp. The resulting sunspot-detection statistics are summarized in Table 4.

The moderate mean confidence reflects the expected domain shift between the space-based HMI continuum images used for model training and the ground-based Hα images used in this external test. Despite differences in wavelength, spatial sampling, seeing conditions, and instrumental response, SOLARIS identifies multiple sunspot structures per frame, with an average detection rate

exceeding five instances per image. These results should therefore be interpreted as an initial demonstration of cross-domain applicability.

4.2. Limitations

Although SOLARIS achieves high performance across the three target classes, some limitations remain that highlight the challenges inherent to multi-feature solar detection using continuum intensitygrams alone.

First, the detection of prominences remains the most difficult task, primarily due to their intrinsically low contrast and strong dependence on limb projection effects. Continuum images do not capture the thermal or emissive structure of prominences as clearly as EUV wavelengths (e.g., 304 Å), which limits the model's ability to distinguish faint or partially occulted structures. This class imbalance in physical visibility likely contributes to the lower mAP@0.5 and mAP50–95 values achieved for prominence detections.

Second, the reliance on continuum intensitygrams restricts the detection of features whose magnetic or thermal signatures are better expressed in magnetograms or EUV channels. Coronal holes, for instance, may appear fragmented or visually ambiguous in continuum data due to their weak contrast relative to the surrounding quiet Sun. Although the model performs well overall, its accuracy is ultimately limited by what can be extracted from the HMI continuum alone.

Finally, the segmentation-derived masks used for morphological refinement are sensitive to bounding-box geometry. While watershed segmentation performs well for large coherent regions, its accuracy decreases for highly irregular shapes or for objects near the solar limb, where foreshortening affects intensity continuity.

However, SOLARIS constitutes a concrete step toward scalable, automated solar activity monitoring. Its combination of high performance, multi-class capability, and efficient inference makes it suitable for integration into long-term monitoring pipelines, educational observatories, and multi-instrument solar-physics workflows.

5. Conclusions

In this work, we presented SOLARIS, a deep-learning framework based on the YOLO26 architecture designed for the automated detection of sunspots, coronal holes, and prominences using SDO/HMI continuum intensitygrams. The system incorporates a physically motivated preprocessing step, a Ray Tune–optimized set of hyperparameters, and a postprocessing pipeline that refines detections through optional segmentation. Together, these components enable a robust and scalable approach to multi-feature solar activity recognition.

The model demonstrates strong performance across all three feature classes. On the test set, SOLARIS achieves an mAP@0.5 of 0.930 (mAP@0.5–0.95 of 0.573), with particularly strong detection performance for sunspots and coronal holes. Prominence detection remains the most challenging task due to intrinsic limb-projection effects and low-contrast morphology in continuum images, but

the model nevertheless performs reliably under a broad range of visibility conditions. Visual inspection confirms that detections are stable for both isolated structures and complex multi-feature scenes.

The stability of the training process, supported by smooth convergence of loss functions and consistent validation metrics, indicates that the optimized hyperparameter configuration effectively balances detection accuracy and model complexity. The system generalizes well across the diverse solar states sampled between May and August 2024, including periods of heightened activity associated with Solar Cycle 25.

While the present implementation focuses on continuum images alone, our results show that meaningful multi-class detection is achievable even in this limited modality. This opens the door to rapid, low-latency characterization of photospheric and chromospheric structures in operational settings. The modular pipeline also enables straightforward integration with multi-wavelength observations, temporal tracking frameworks, and larger-scale solar-monitoring pipelines.

As a practical extension of this work, the SOLARIS framework is now being incorporated into the operational workflow of the LANCE–UANL Solar Observatory, where it will support continuous monitoring of photospheric activity from ground-based observations. This deployment provides an opportunity to evaluate the system under real observational constraints—such as atmospheric variability, instrumental noise, and daily acquisition gaps—offering valuable feedback for improving the robustness of the model. Its integration into the laboratory's monitoring infrastructure also enables the development of automated tracking tools and real-time alerts, laying the groundwork for future forecasting experiments and long-term solar-activity studies.

Acknowledgements M. Rodríguez-Martínez acknowledges the support of the DGAPA PAPIIT project IN115423 and the SECIHTI Frontier Science 2023 grant CF-2023-G-364. He also thanks J. Mejía-Orozco, M. Peralta-Mondragón, A. Varela-Mendez, C. A. Andrade-Pineda, and R. Gutiérrez-Zalapa for their contributions to image acquisition. This work was co-funded by the European Union under Project 101183150 – OCEANS.

Data & Repository Information

All results, optimization logs, and analysis files used in this study are available in the public repository Jara Reyna et al. (2026).

References

Gong, L., Yang, Y., Feng, S., Dai, W., Liang, B., Xiong, J.: 2024, Solar Active Regions Detection and Tracking Based on Deep Learning. *Sol. Phys.* **299**, 121. DOI. ADS.

Hathaway, D.H.: 2015, The Solar Cycle. *Living Reviews in Solar Physics* **12**. DOI.

Jara Reyna, M.A., Guerrero Peña, C.A., Rodriguez Vazquez, A.F., Romero-Hernandez, E., Tapia, J.J., Martínez, D., Ibarra, Y., Ayala, S., Rodríguez Martínez, M.: 2026, YOLO26 Object Detection Model — Reproducible Training, Evaluation & Results Repository. *Zenodo Repository*. DOI. https://doi.org/10.5281/zenodo.19531480.

Jocher, G., Qiu, J.: 2026, *Ultralytics YOLO26*. https://github.com/ultralytics/ultralytics.

Kauristie, K., Andries, J., Beck, P., Berdermann, J., Berghmans, D., Cesaroni, C., De Donder, E., de Patoul, J., Dierckxsens, M., Haralambous, H., et al.: 2021, Space Weather Services for Civil Aviation–Challenges and Solutions. *Remote Sensing* **13**, 3685. DOI.

Ladegaard, A., Wingham, M., Heron, S.: 2017, *The Solar Limb Darkening Toolkit (SLDTk)*, https://github.com/LandingEllipse/SLDTk.

Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., Talwalkar, A.: 2018, Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *Journal of Machine Learning Research* **18**, 1. http://jmlr.org/papers/v18/16-558.html.

Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J.E., Stoica, I.: 2018, Tune: A Research Platform for Distributed Model Selection and Training. *arXiv e-prints* **arXiv:1807.05118**. DOI.

Lowder, C., Qiu, J., Leamon, R.: 2017, Coronal Holes and Open Magnetic Flux over Cycles 23 and 24. *Sol. Phys.* **292**, 18. DOI. ADS.

Moon, B., Jeong, D.-G., Oh, S., Sohn, J.: 2017, Variation in Solar Limb Darkening Coefficient Estimated from Solar Images Taken by SOHO and SDO. *Journal of Astronomy and Space Sciences* **34**, 99. DOI. https://doi.org/10.5140/JASS.2017.34.2.99.

Müller, D., Nicula, B., Felix, S., Verstringe, F., Bourgoignie, B., Csillaghy, A., Berghmans, D., Jiggens, P., García-Ortiz, J.P., Ireland, J., Zahniy, S., Fleck, B.: 2017, JHelioviewer. Time-dependent 3D visualisation of solar and heliospheric data. *Astron. Astrophys.* **606**, A10. DOI. ADS.

Opgenoorth, H.J., Wimmer-Schweingruber, R.F., Belehaki, A., Berghmans, D., Hapgood, M., Hesse, M., Kauristie, K., Lester, M., Lilensten, J., Messerotti, M., Temmer, M.: 2019, Assessment and recommendations for a consolidated European approach to space weather – as part of a global space weather effort. *Journal of Space Weather and Space Climate* **9**, A37. DOI.

Pan, J., Liu, J., Fang, S., Liu, R.: 2025, SARD: A YOLOv8-Based System for Solar Active Region Detection with SDO/HMI Magnetograms. *Sol. Phys.* **300**, 111. DOI. ADS.

Parenti, S.: 2014, Solar Prominences: Observations. *Living Reviews in Solar Physics* **11**, 1. DOI. ADS.

Pesnell, W.D., Thompson, B.J., Chamberlin, P.C.: 2012, The Solar Dynamics Observatory (SDO). *Sol. Phys.* **275**, 3. DOI. ADS.

Petrovay, K.: 2020, Solar cycle prediction. *Living Reviews in Solar Physics* **17**, 2. DOI.

Pulkkinen, T.: 2007, Space Weather: Terrestrial Perspective. *Living Reviews in Solar Physics* **4**. DOI.

Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: 2019, Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 658. DOI.

Sandford, E., Auchère, F., Mortier, A., Hayes, L.A., Müller, D.: 2025, Presenting 28 years of Sun-as-a-star extreme ultraviolet light curves from SOHO EIT. *Astronomy and Astrophysics* **699**, A353. DOI.

Schrijver, C.J.: 2007, A Characteristic Magnetic Field Pattern Associated with All Major Solar Flares and Its Use in Flare Forecasting. *Astrophysical Journal Letters* **655**, L117. DOI.

Schrijver, C.J.: 2015, Socio-Economic Hazards and Impacts of Space Weather: The Important Range Between Mild and Extreme. *Space Weather* **13**, 524. DOI.

Thompson, W.T.: 2006, Coordinate systems for solar image data. *Astron. Astrophys.* **449**, 791. DOI. ADS.

Toriumi, S., Wang, H.: 2019, Flare-productive active regions. *Living Reviews in Solar Physics* **16**. DOI.

Wiegelmann, T., Thalmann, J.K., Solanki, S.K.: 2014, The magnetic field in the solar atmosphere. *Astronomy and Astrophysics Review* **22**, 78. DOI.