

A Hybrid CNN–Transformer Network to Enhance Solar Magnetogram Resolution for Flare Predictive Analytics

Vishakha Singh

National Institute of Technology Delhi

Divya Punia

dpunia@nitdelhi.ac.in

National Institute of Technology Delhi

V.S. Pandey

National Institute of Technology Delhi

Ajay K Sharma

National Institute of Technology Delhi

Research Article

Keywords: Solar Flares, Forecasting, Magnetic Fields, Photosphere; Active Regions, Magnetic Fields, Image Processing, Data Analysis, Deep learning, Magnetograms, Super Resolution, Resolution Enhancement

Posted Date: March 26th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9159480/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

A Hybrid CNN–Transformer Network to Enhance Solar Magnetogram Resolution for Flare Predictive Analytics

Vishakha¹ · Divya Punia^{2,*} · V. S. Pandey¹ · Ajay K. Sharma³

© The author(s) • • • •

Abstract : The disparity in spatial resolution between the SOHO/MDI and SDO/HMI magnetograms creates inconsistencies that hinder reliable long term studies of solar magnetic fields and reduce flare forecasting. To solve this problem, the MagRes-Net a hybrid super-resolution Network that uses convolutional neural networks (CNNs) to extract features locally and transformer-based self attention to extract global gradients. The designed network reconstructs high-resolution magnetograms from lower resolution images while maintaining physical structures realistic features of the magnetic field. The model is trained using MDI-HMI co-aligned image pairs and implemented physics-aware constraints into the model, such as gradient consistency, and magnetic flux conservation. Through quantitative evaluations, MagRes-Net was better than interpolation and CNN-based methods in terms of peak signal-to-noise ratio (PSNR), mean squared error (RMSE), Learned Perceptual Image Patch Similarity (LPIPS) and structural similarity index measure (SSIM). Quantitatively, AR coverage increases from 12.23% in native MDI observations to 13.25% in the super-resolved output, approaching the SDO/HMI reference value of 14.48%. it indicates the model successfully recovers resolution-dependent attributes. Additionally, spectral, noise robustness and perceptual analyses demonstrated that MagRes-Net preserves structure at all frequencies, and thus can improve resolution dependent magnetic diagnostic techniques, including; continuity of polarity inversion lines, morphology of strong field regions, and magnetic flux distribution. Therefore, these results demonstrate that the proposed method bridges the cross-instrument resolution disparities, allowing for the creation of homogeneous magnetogram data sets that are capable for long term studies relevant for solar activity and improved flare forecasting.

Keywords: Solar Flares, Forecasting; Magnetic Fields, Photosphere; Active Regions, Magnetic Fields; Image Processing; Data Analysis; Deep learning; Magnetograms; Super Resolution; Resolution Enhancement

Article Highlights

- Hybrid CNN–Transformer super-resolution reduces the MDI–HMI domain gap while preserving magnetic topology.
- Physics-aware diagnostics show improved flux consistency and PIL continuity approaching the HMI reference.
- Spectral and robustness tests indicate stable frequency-domain fidelity under noise perturbations.

1. Introduction

Solar eruptions phenomena are a sudden release of the magnetic energy accumulated on the Sun's surface in areas that are typically referred to as Active regions. Solar flares are rapid emissions of magnetic energy, since these eruptions can have an impact on technology on earth, such as satellites and power

¹ Department of Applied Sciences, National Institute of Technology Delhi, New Delhi, India

vishakha@nitdelhi.ac.in

² Department of Electronics and Communication Engineering, National Institute of Technology Delhi, New Delhi, India

dpunia@nitdelhi.ac.in (Corresponding Author)

³ Department of Computer Science and Engineering, National Institute of Technology Jalandhar, Jalandhar, India

grids, there is a need for accurate prediction. However, it is difficult to measure the coronal magnetic fields directly. Thus, photospheric magnetograms serve as the best proxy to determine the magnetic complexity at the base of the corona and how it evolves before a major eruption occurs. (Schrijver 2007; Bobra and Couvidat 2015). Numerous studies have demonstrated that flare-productive active regions are statistically associated with strong-field concentrations, large unsigned magnetic flux, and extended polarity inversion lines (PILs) exhibiting steep magnetic-field gradients, which act as proxies for magnetic shear and non-potentiality. Long-term investigations of solar magnetic variability and flare occurrence require magnetogram observations that remain consistent across multiple instruments and solar cycles. The SOHO/MDI instrument provides an invaluable pre-*SDO* archive, whereas SDO/HMI defines the modern reference standard for high-resolution magnetic-field measurements (Scherrer et al. 1995, 2012). However, the substantial spatial-resolution difference between MDI and HMI introduces a practical gap between instruments. Single-image super-resolution (SR) offers a data-driven approach for mitigating this mismatch by learning a mapping from low-resolution to high-resolution observations. Deep CNN-based SR initiated with SRCNN (Dong et al. 2016) demonstrated strong improvements, with residual and attention-based models further enhancing performance (Lim et al. 2017; Zhang et al. 2018).

Transformers provide a complementary mechanism by modelling nonlocal dependencies. Vision Transformers capture long-range context (Dosovitskiy et al. 2021), while Swin Transformers improve efficiency via shifted windows (Liu et al. 2021). These ideas were adapted to SR through SwinIR (Liang et al. 2021). Unlike conventional SR approaches focused primarily on visual reconstruction, this study treats magnetogram enhancement as an instrument-bridging problem motivated by space-weather applications.

2. Related Work

2.1. Learning-based SR for solar magnetograms

Physics-aware SR considerations were emphasised in early MDI/HMI bridging studies (Jungbluth et al. 2019). SR was also applied via degradation-and-reconstruction strategies (Rahman et al. 2020). Gradient-guided and structure-preserving strategies were later introduced for MDI→HMI enhancement (Dou et al. 2022), while attention-aided CNNs improved fidelity on active-region magnetograms (Xu et al. 2024). Full-disk enhancement has also been explored (Qin et al. 2024). Recent physically motivated efforts also explicitly target both super-resolution and cross-instrument calibration to enable homogeneous long-baseline magnetogram diagnostics across observing eras (Munoz-Jaramillo et al. 2024). More recently, diffusion-style SR emerged as an uncertainty-aware alternative (Ramunno et al. 2025). Related to the instrument-bridging objective, unpaired instrument-to-instrument translation has also been demonstrated as a pathway to homogenize multi-instrument solar observations without requiring strict co-temporality (Jarolim et al. 2025).

Deep learning has also been widely explored for enhancing and denoising solar observations beyond magnetograms, demonstrating that learned priors can recover fine-scale structure while suppressing instrumental noise and degradations (Díaz Baso and Asensio Ramos 2018; Asensio Ramos and Díaz Baso 2019).

2.2. CNN and ViT-Transformer SR backbones

CNN-based SR began with SRCNN (Dong et al. 2016), followed by deeper residual networks such as EDSR (Lim et al. 2017). Channel attention was introduced through RCAN (Zhang et al. 2018). Global context modelling via Transformers was introduced through ViT (Dosovitskiy et al. 2021) and Swin Transformer (Liu et al. 2021), later adapted to SR via SwinIR (Liang et al. 2021).

As seen in Table 1, recent methods for super-resolving solar magnetograms follow one of two main directions: local-detail enhancement or probabilistic reconstruction. Local fidelity improvement techniques based on CNNs often lack long-range spatial consistency, whereas diffusion-based techniques can model uncertainty but come with greater computational cost and reduced interpretability. Physically motivated calibration frameworks exist, but they are not optimized for high-fidelity super-resolution. The proposed MagRes-Net framework addresses this gap by modelling both local and global magnetic structures while incorporating physics-aware constraints for more meaningful magnetogram reconstruction.

Table 1. Comparison of representative prior super-resolution approaches for solar magnetograms and the proposed MagRes-Net framework. The table highlights methodological differences, physical constraints, and key limitations.

(a) Representative prior state-of-the-art super-resolution approaches for solar magnetograms				
Work	Target	Method	Physics	Limitation
(Díaz Baso and Asensio Ramos 2018)	HMI (synthetic)	CNN SR	PSF-aware	Synthetic bias
(Jungbluth et al. 2019)	MDI→HMI	Physics SR	Explicit constraints	Low capacity
(Rahman et al. 2020)	HMI	CNN+GAN	Indirect validation	Weak physics
(Dou et al. 2022)	MDI→HMI	Gradient SR	Polarity aware	Local-only
(Xu et al. 2024)	AR	Attention CNN	Strong-field	Patch-based
(Qin et al. 2024)	Full disk	Regression SR	Correlation metrics	Over-smoothing
(Munoz-Jaramillo et al. 2024)	Cross-inst.	ML calibration	Strong physics	Not SR-focused
(Ramunno et al. 2025)	Magnetograms	Diffusion SR	Uncertainty	Expensive
(b) Proposed MagRes-Net Model				
Aspect	Architecture	Key Strength	Physics Awareness	Current Limitation
MagRes-Net	Hybrid CNN–Transformer SR with attention-guided feature fusion	Targets MDI→HMI and Captures local gradients and global magnetic topology	Flux preservation, gradient consistency, PIL structure, spectral fidelity	Temporal validation remains to be investigated.

3. Data Set and Preprocessing

Motivated by these developments, the present work proposes **MagRes-Net**, a hybrid architecture combining CNN locality bias with Transformer global receptive fields tailored to solar magnetograms.




Figure 1. Co-temporal line-of-sight (LOS) magnetogram submaps of NOAA AR 11121 (06 November 2010), illustrating the resolution gap between SOHO/MDI and SDO/HMI observations. Left: lower-resolution MDI input, which under-resolves compact flux elements and sharp polarity transitions. Right: higher-resolution HMI reference used for super-resolution training and evaluation.




Figure 2. Preprocessing pipeline used to construct paired low- and high-resolution magnetogram samples. The workflow includes temporal pairing, WCS-based co-alignment, active-region ROI extraction, clipping, and normalisation prior to super-resolution training.

Co-temporal SOHO/MDI and SDO/HMI LOS magnetograms were paired within a ± 8 min temporal window to minimise discrepancies caused by short-term solar evolution. A representative example of the native resolution gap between the two instruments is shown in Fig. 1, while the overall preprocessing workflow used to generate paired training samples is summarized in Fig. 2. Observations with heliographic longitudes $|\lambda| > 60^\circ$ were excluded to reduce projection effects near the solar limb, and corrupted or incomplete records were removed during quality control.

Spatial co-alignment between MDI and HMI magnetograms was performed using World Coordinate System (WCS) metadata to ensure pixel-level correspondence between low-resolution inputs and high-resolution references (Thompson 2006). Active-region cutouts were extracted from the aligned maps, resulting in 16,574 co-registered samples of size 128×128 pixels. After visual and statistical quality screening, 8,287 high-quality pairs were retained, from which 5,232 activity-rich submaps were selected for model training and evaluation.

The magnetogram values were clipped to a value of ± 3000 G to remove extreme outliers while preserving strong-field regions. The clipped data were then normalized to the range $[-1, 1]$ to stabilize network training. After reconstruction, outputs were de-normalized back to physical units (Gauss) to enable direct quantitative comparison.

The split into train/validation/test was performed at the active-region level rather than at the patch level. Therefore, no overlap occurs between active regions used for training, validation, and testing, providing a more reliable estimate of model generalization.

AR 11121 on 06 November 2010 was selected as the case study because it provides a suitable test-bed for evaluating the preservation of compact flux concentrations, extended PIL morphology, and flare-relevant gradient structure.

4. Methodology

The MagRes-Net architecture is designed to reconstruct high-resolution solar magnetograms from low-resolution observations of the same regions while preserving the underlying physical structure of magnetic systems. Solar magnetograms exhibit features on many spatial scales across the solar surface, so an effective super-resolution method must incorporate both fine-scale local gradient information and large-scale global structural context. To meet these demands, MagRes-Net integrates convolutional feature extraction with global context modeling through a Transformer-based framework, forming a hybrid model.

The complete pipeline consists of five main stages: shallow feature extraction, Transformer-based global representation learning, convolutional local feature refinement, attention-guided feature fusion, and final reconstruction through pixel-shuffle upsampling. The overall architecture is illustrated in Fig. 3.




Figure 3. Overall architecture of MagRes-Net. The network processes a low-resolution MDI magnetogram through shallow feature extraction followed by two parallel branches. Their features are fused using an attention-guided Feature Fusion Block and subsequently refined through reconstruction convolution and pixel-shuffle upsampling to generate the super-resolved magnetogram.

4.1. Shallow Feature Extraction

Given an input low-resolution magnetogram $I_{LR} \in \mathbb{R}^{H \times W}$, the network first extracts shallow spatial features using a 3×3 convolutional layer:

$$F_0 = \text{Conv}_{3 \times 3}(I_{LR}) \quad (1)$$

The first step in generating a feature map F_0 , is to preserve the spatial resolution of the input with increased dimensions. Shallow feature extraction layers capture primary edge features and magnetic gradient information necessary for subsequent reconstruction processes. By using a small receptive field at this point of processing; the local magnetic structure will be retained prior to applying higher level feature transformations.

4.2. Transformer Global-Context Branch

The relationships among magnetic field configurations of active region of the sun in terms of both scale and topological relationship with surrounding larger-scale structures depend on a large range of spatial relationships. Because convolutional operations typically cannot capture these longer-range spatial relationships, a Transformer-based branch is incorporated into MagRes-Net to capture the contextual relationships at large scale among spatially distributed tokens based upon their positions in the embedded token space that has been produced by mapping the initial shallow feature map F_0 as shown below:

$$T_0 = \text{Embed}(F_0) \quad (2)$$

Each token represents a local patch of the magnetogram. Tokens are processed using Transformer blocks composed of multi head self attention (MHSA) and feed forward multilayer perceptron layers with residual connections and normalization layers. The MHSA mechanism computes pairwise relationships between tokens:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (3)$$

where Q , K , and V denote query, key, and value matrices derived from token embeddings. The MHSA enables the network to model global interactions between spatially distant magnetic regions, which is especially useful in capturing coherent active region morphology. The Transformer based representations have demonstrated high performance in modeling long range dependencies for image restoration and super resolution tasks (Dosovitskiy et al. 2021; Liu et al. 2021).

4.3. CNN Local-Gradient Branch

While Transformers do represent the global dependency, there are additional requirements to accurately reconstruct the local magnetic gradient and the boundary for sharp polarity. The parallel convolutional branch uses stacked convolutional blocks to process the shallow feature. As such, each of the blocks include a 3×3 convolution; nonlinear activation; and normalization, which will be described below. To bring together the local and global representation, the fusion strategy illustrated in Fig. 4 was employed.

$$F_{i+1} = \sigma(\text{BN}(\text{Conv}_{3 \times 3}(F_i))) \quad (4)$$

These layers improve on local structure details in each layer so that the network can find where to expect finer scale flux concentrations and polarity reversal lines. Convolutional operation is well suited to model local spatial continuity and gradients based upon local features; these are two of the primary characteristics of solar magnetogram data.




Figure 4. Feature Fusion Block used to integrate CNN and Transformer features. After feature concatenation, a 1×1 convolution aligns channel representations, followed by Efficient Channel Attention (ECA) for adaptive channel weighting. A 3×3 convolution further refines the fused representation, while residual connections preserve original feature information.

4.4. Feature Fusion Block

The outputs from the CNN and Transformer branches are combined using the Feature Fusion Block (Figure 4) after they have been independently extracted as features. The purpose of this block is to merge representations of complementary information from each branch while suppressing redundant or noisy activations. The CNN feature F_{cnn} , and the Transformer feature F_{tr} are first concatenated along channels:

$$F_{\text{cat}} = [F_{\text{cnn}}; F_{\text{tr}}] \quad (5)$$

To reduce dimensional redundancy and align channel representations, a 1×1 convolution is applied:

$$a = 1 \quad (6)$$

$$F_{\text{proj}} = \text{Conv}_{a \times a}(F_{\text{cat}}) \quad (7)$$

This projection step performs channel mixing and dimensionality reduction while preserving spatial resolution.

4.5. Efficient Channel Attention

The use of Channel Attention Mechanisms enables the network to focus on the relevant information, while minimizing the impact of non-relevant information. The use of an Efficient Channel Attention (ECA), from Wang et al. (2020), (Wang et al. 2020) for MagRes-Net, allows ECA to provide channel-wise attention without a need for a dimensionality reduction. Figure 5, shows how the global average pooling method first gathers the spatial information:

$$z_c = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W F_c(i, j) \quad (8)$$

A lightweight 1D convolution models cross-channel interactions to generate attention weights:

$$w = \sigma\left(\text{Conv}^{aD}(z)\right), \quad a = 1 \quad (9)$$

These weights are then applied to rescale the feature channels:

$$F_{\text{eca}} = w \odot F_{\text{proj}} \quad (10)$$

The network can be able to give greater priority to magnetically relevant aspects such as strong-field areas and polarity boundaries while reducing noise.




Figure 5. Efficient Channel Attention (ECA) mechanism used in the fusion module. Global average pooling summarizes channel statistics, which are processed through a lightweight convolution to generate adaptive channel-wise scaling weights.

Algorithm 1 Multi-scale smoothing for robust PIL extraction

Require: LOS magnetogram M (FITS); ROI corners r_1, r_2 ; thresholds τ_+, τ_- ; Gaussian scales $\mathcal{S} = \{s_1, s_2, s_3\}$; gradient floor $g_{\min}$; minimum contour length $L_{\min}$

Ensure: Binary PIL mask Π

- 1: $B \leftarrow \text{CropWCS}(M, r_1, r_2)$
 - 2: $\Omega_+ \leftarrow [B > \tau_+]$; $\Omega_0 \leftarrow [\tau_- \leq B \leq \tau_+]$; $\Omega_- \leftarrow [B < \tau_-]$
 - 3: **for all** $s \in \mathcal{S}$ **do**
 - 4: $F_s \leftarrow \text{Gauss}(B, s)$
 - 5: $\tilde{B} \leftarrow \Omega_+ \odot F_{s_1} + \Omega_0 \odot F_{s_2} + \Omega_- \odot F_{s_3}$
 - 6: $\Pi_0 \leftarrow \text{ZeroContour}(\tilde{B})$
 - 7: $G \leftarrow \|\nabla \tilde{B}\|$
 - 8: $\Pi_1 \leftarrow \Pi_0 \wedge [G > g_{\min}]$
 - 9: $\Pi \leftarrow \text{RemoveSmall}(\Pi_1, L_{\min})$
 - 10: **return** Π
-

4.6. Robust PIL Extraction

The PIL extraction algorithm is shown in Algorithm 1, the PILs show areas where the magnetic field (Bz) has changed signs. They are often used as indicators for complex magnetic structures and flare activity. After an active region area is cropped from the magnetogram data that was reconstructed from WCS coordinates, the area is separated into areas with positive, neutral and negative polarity. The segmentation map will be filtered using a multilevel Gaussian filter to remove noise at the pixel scale while retaining the large-scale structures. Once the filtering is done, a combined magnetic field map can be made. The zero level contours of this combined map can now be used to define candidate inversion lines. Weak gradient crossings may not have physical polarity transitions, the gradient magnitude of the composite field is computed and weak-gradient crossings are rejected. Finally, short isolated contour fragments are removed using a minimum-length threshold, producing a binary PIL map suitable for quantitative analysis of flare-relevant properties such as PIL length and gradient-weighted complexity.

Algorithm 1 summarizes the robust procedure adopted to extract polarity inversion lines (PILs) from reconstructed magnetograms.

4.7. Reconstruction and Upsampling

After feature fusion and refinement, the aggregated representation is passed through a reconstruction convolution layer followed by a pixel-shuffle upsampling module. Pixel shuffle rearranges feature channels into spatial resolution, enabling efficient sub-pixel convolution-based upscaling (Shi et al. 2016). The final high-resolution magnetogram $\hat{I}_{HR}$ is obtained as

$$\hat{I}_{HR} = \text{PixelShuffle}(\text{Conv}_{3 \times 3}(F_{cca})) \quad (11)$$

Table 2. Hyperparameters for training MAGRES-NET.

Setting	Value
Patch size	128×128 ; upsampling $s = 4$
Data split	4000 / 616 / 616 (train/val/test)
Optimizer	AdamW; LR 2×10^{-4} (cosine); WD 10^{-4}
Training	Batch 16; 100 epochs
CNN branch	8 blocks; 64 channels
ViT branch	Embed 96; 4 heads; 4 blocks
Upsampling	Pixel-shuffle

4.8. Training Objective and Physical Constraints

Precise super-resolution of photospheric magnetograms requires both pixel-level reconstruction accuracy and preservation of physically meaningful magnetic structure. To achieve this, the training objective combines multiple loss components. The Hyperparameters for training is mentioned in Table 2.

The primary reconstruction loss is the Charbonnier loss:

$$\mathcal{L}_{\text{char}} = \frac{1}{N} \sum_{i=1}^N \sqrt{\left(\hat{I}_{HR}^{(i)} - I_{HR}^{(i)}\right)^2 + \epsilon^2} \quad (12)$$

To encourage structural fidelity, a gradient-consistency loss is added:

$$\mathcal{L}_{\nabla} = \|\nabla \hat{I}_{HR} - \nabla I_{HR}\|_1 \quad (13)$$

Additionally, a magnetic-flux conservation term is introduced to penalize deviations in the total signed magnetic flux:

$$\mathcal{L}_{\Phi} = |\Phi(\hat{I}_{HR}) - \Phi(I_{HR})| \quad (14)$$

The total optimization objective is therefore

$$\mathcal{L} = \mathcal{L}_{\text{char}} + \lambda_{\nabla} \mathcal{L}_{\nabla} + \lambda_{\Phi} \mathcal{L}_{\Phi} \quad (15)$$

4.9. Training Details

All model variants were trained for 100 epochs using the AdamW optimizer with learning rate 2×10^{-4} , $\beta_1 = 0.9$, $\beta_2 = 0.999$, and weight decay 10^{-4} . The learning rate was scheduled using cosine decay with a short warm-up phase. The weighting parameters were set to $\lambda_{\nabla} = 0.10$ and $\lambda_{\Phi} = 0.05$ in all experiments.

This training strategy enables the model to balance pixel-wise reconstruction accuracy with preservation of physically meaningful magnetic structures.

5. Evaluation Metrics

5.1. Image-level reconstruction metrics

Let $\hat{X}$ denote the super-resolved magnetogram and X the co-registered HMI reference over N pixels. Pixel-level reconstruction accuracy is quantified using the mean squared error (MSE) and root mean squared error (RMSE),

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (\hat{X}_i - X_i)^2, \quad \text{RMSE} = \sqrt{\text{MSE}}. \quad (16)$$

The peak signal-to-noise ratio (PSNR) is defined as

$$\text{PSNR} = 10 \log_{10} \left(\frac{X_{\max}^2}{\text{MSE}} \right). \quad (17)$$

Structural similarity is quantified using SSIM,

$$\text{SSIM}(X, \hat{X}) = \frac{(2\mu_X \mu_{\hat{X}} + C_1)(2\sigma_{X\hat{X}} + C_2)}{(\mu_X^2 + \mu_{\hat{X}}^2 + C_1)(\sigma_X^2 + \sigma_{\hat{X}}^2 + C_2)}. \quad (18)$$

Perceptual similarity is evaluated using LPIPS,

$$\text{LPIPS} = \sum_l \frac{1}{H_l W_l} \sum_{h,w} \left\| w_l \odot (\phi_l(X) - \phi_l(\hat{X})) \right\|_2^2. \quad (19)$$

Although Learned Perceptual Image Patch Similarity (LPIPS) was originally designed for natural-image perception, it has recently been applied to scientific image reconstruction tasks as a complementary measure of structural similarity. LPIPS evaluates differences in deep feature representations extracted from a pretrained neural network. In the context of solar magnetograms, LPIPS provides an additional assessment of whether reconstructed maps preserve coherent structural patterns and polarity transitions beyond simple pixel-wise similarity metrics.

Frequency-domain consistency is evaluated through Fourier analysis. Let $\mathcal{F}(X)$ denote the 2D Fourier transform, with magnitude and phase spectra

$$M_X = |\mathcal{F}(X)|, \quad \Phi_X = \arg(\mathcal{F}(X)). \quad (20)$$

5.2. Additional spectral and perceptual metrics

To further evaluate structural and spectral fidelity, we compute the Spectral Angle Mapper (SAM), frequency-domain correlation, and phase error statistics.

5.2.0.1. Spectral Angle Mapper (SAM). SAM measures the angular similarity between the reconstructed and reference magnetograms. For vectors X and $\hat{X}$, it is defined as

$$\text{SAM}(X, \hat{X}) = \arccos \left(\frac{\sum_{i=1}^N X_i \hat{X}_i}{\sqrt{\sum_{i=1}^N X_i^2} \sqrt{\sum_{i=1}^N \hat{X}_i^2}} \right). \quad (21)$$

Lower SAM values indicate higher similarity.

5.2.0.2. Frequency-domain correlation. To assess spectral consistency, we compute the Pearson correlation between the Fourier magnitude spectra of the reference and reconstructed magnetograms. Let $M_X = |\mathcal{F}(X)|$ and $M_{\hat{X}} = |\mathcal{F}(\hat{X})|$. The frequency correlation is

$$\rho_f = \frac{\sum (M_X - \bar{M}_X)(M_{\hat{X}} - \bar{M}_{\hat{X}})}{\sqrt{\sum (M_X - \bar{M}_X)^2} \sqrt{\sum (M_{\hat{X}} - \bar{M}_{\hat{X}})^2}}. \quad (22)$$

Values close to unity indicate strong preservation of spectral energy.

5.2.0.3. Phase error statistics. Fourier phase contains spatial structural information. The phase difference between reconstructed and reference images is computed as

$$\Delta\Phi = \text{wrap}(\Phi_{\hat{X}} - \Phi_X), \quad (23)$$

where the wrapping operation constrains the phase difference to

$$\Delta\Phi \in [-\pi, \pi]. \quad (24)$$

The mean and standard deviation of $\Delta\Phi$ quantify phase consistency across spatial frequencies.

6. Results and Discussion

6.1. Visual and quantitative reconstruction performance

Quantitative reconstruction performance on the held-out test set is summarized in Table 3, while qualitative visual comparisons for NOAA AR 11121 are presented in Fig. 7. Among the evaluated approaches, MagRes-Net attains the highest reconstruction fidelity, improving PSNR by 7.17 dB over bicubic interpolation and by 3.51 dB over the Deep CNN baseline, while reducing RMSE by 56.9% and 43.3%, respectively.

6.2. Physical validation: flux preservation, AR structure, and PIL topology

High pixel-level fidelity is necessary but not sufficient for magnetogram super-resolution intended for space-weather use, because flare-relevant predictors depend on physically meaningful magnetic structure

Table 3. Image-level SR metrics on the held-out test set. Higher PSNR is better; lower RMSE is better. RMSE values are reported in Gauss.

Method	PSNR (dB)	RMSE (G)
Bicubic	37.136	6.412
Deep CNN	40.804	4.873
MagRes-Net (CNN+ViT)	44.310	2.76

rather than visual similarity alone. We therefore validate the reconstructed fields using flux-based diagnostics, threshold-defined active-region structure, and polarity inversion line topology, which are widely used as proxies for magnetic complexity and flare productivity (Schrijver 2007; Bobra and Couvidat 2015; Hoeksema et al. 2014).

Flux preservation. Figures 6 and 9 compare signed and unsigned flux distributions derived from native MDI, the MagRes-Net output, and the HMI reference for NOAA AR 11121.

Active-region structure and strong-field morphology. We apply a strong-field threshold of $|B_z| > 100$ G and compare the resulting contours (Fig. 10). The corresponding μ -map representation with extracted active-region boundaries is shown in Fig. 11. The MagRes-Net output yields more continuous boundaries and recovers additional strong-field structure that agrees better with HMI. This improvement is also reflected in the AR coverage values (Table 5).

PIL topology and gradient-linked complexity. Because flare productivity is strongly associated with the presence of long, coherent PILs embedded in high-gradient magnetic environments (Schrijver 2007; Georgoulis and Rust 2007), we examine whether super-resolution improves PIL continuity without creating spurious reversals (Fig. 13).




Figure 6. Signed magnetic flux comparison for NOAA AR 11121. From left to right: native MDI, MagRes-Net super-resolved output, and HMI reference.

Table 4. Quantitative comparison of the proposed MagRes-Net model with representative state-of-the-art super-resolution approaches for solar magnetograms. Higher PSNR and SSIM indicate better reconstruction quality, while lower LPIPS and RMSE indicate improved perceptual fidelity and reduced reconstruction error.

Model	PSNR $\uparrow$ (dB)	SSIM $\uparrow$	LPIPS $\downarrow$	RMSE $\downarrow$ (G)
SPSR ((Dou et al. 2022))	38.99	0.89	0.33	21.10
LDM ((Ramunno et al. 2025))	39.60	0.79	0.50	31.16
SolarCNN ((Xu et al. 2024))	40.12	0.67	0.62	41.47
MagRes-Net (Ours)	44.31	0.97	0.06	2.76

To quantify the performance of the proposed method relative to , we evaluated it against three representative existing super-resolution approaches for solar magnetograms: the structure-preserving super-resolution method (SPSR) proposed by (Dou et al. 2022), the latent diffusion model (LDM) introduced by (Ramunno et al. 2025), and the attention-aided convolutional network (SolarCNN) developed by (Xu et al. 2024). The comparison is including PSNR, SSIM, LPIPS, and RMSE. As summarized in Table 3, the proposed MagRes-Net model consistently outperforms the baseline approaches across the




Figure 7. Visual comparison for NOAA AR 11121 across reconstruction methods. Native MDI under-resolves compact flux elements and narrow polarity inversion lines. Bicubic interpolation and the Deep CNN baseline partially recover structure but remain overly smooth. The MagRes-Net reconstruction restores fine-scale magnetic features and polarity inversion line (PIL) continuity, closely matching the HMI reference.

evaluated metrics. Specifically in As summarized in Table 4, the proposed MagRes-Net model consistently outperforms existing state-of-the-art approaches across all evaluation metrics. In particular, MagRes-Net achieves the highest PSNR of 44.31 dB and SSIM of 0.97, indicating superior reconstruction fidelity and structural similarity relative to the HMI reference magnetograms. In comparison, SPSR (Dou et al. 2022) attains a PSNR of 38.99 dB and SSIM of 0.89, while LDM (Ramunno et al. 2025) and SolarCNN (Xu et al. 2024) achieve PSNR values of 39.60 dB and 40.12 dB with SSIM values of 0.79 and 0.67, respectively. the LPIPS value for MagRes-Net is 0.06, which is lower than SPSR (0.33), LDM (0.50), and SolarCNN (0.62), indicating the model preserves fine-scale structures more effectively. Similarly, the RMSE for MagRes-Net has been decreased to 2.76 G in comparison to SPSR (21.10 G) and LDM (31.16 G) as well as SolarCNN (41.47 G). The results show that MagRes-Net uses a hybrid CNN-Transformer architecture and therefore demonstrates high-quality reconstruction of structures using the magnetic field characteristics relative to other SR approaches.

Table 5. Active-region coverage estimated from native MDI, the super-resolved output, and the HMI reference magnetograms.

Data Product	AR Coverage (%)
MDI (Original)	12.23
SR Output (MagRes-Net)	13.25
HMI (Reference)	14.48

6.3. Spectral fidelity, phase consistency, and robustness

To determine if MagRes-Net retains not only spatial-domain similarity but the multi-scale structure as well of solar magnetic field structures, we will measure reconstruction quality in the Fourier domain and through controlled noise additions. In addition to a comparison of the log-magnitude spectra of the HMI reference and super-resolved output (see Fig. 15), an additional frequency-domain error image is




Figure 8. Qualitative comparison with recent super-resolution approaches for solar magnetograms. (a) Structure-preserving super-resolution method of Dou et al. (Dou et al. 2022). (b) Attention-guided SolarCNN model of Xu et al. (Xu et al. 2024). (c) Latent diffusion-based super-resolution framework proposed by Ramunno et al. (Ramunno et al. 2025). (d) Reconstruction generated by the proposed MagRes-Net model (ours). Compared with existing approaches, the proposed method better restores fine-scale magnetic structures and preserves sharper polarity inversion boundaries.




Figure 9. Unsigned magnetic flux comparison for NOAA AR 1121. From left to right: native MDI, MagRes-Net output, and HMI reference.




Figure 10. Comparison of LOS magnetograms thresholded at $|B_z| > 100$ G for NOAA AR 11121. From left to right: native MDI, MagRes-Net output, and HMI reference.




Figure 11. μ -maps ($\mu = \cos \theta$) with active-region boundaries for NOAA AR 11121. Red contours show AR boundaries extracted from the corresponding magnetograms.




Figure 12. Magnetic stress proxy computed from the super-resolved magnetogram, defined as $|\nabla B_z|^2$.

Table 6. Polarity inversion line statistics for NOAA AR 11121.

Metric	MDI (Original)	MagRes-Net	HMI (Reference)
Raw PIL candidates (pixels)	4128	4365	4521
Final PIL length (pixels)	72	84	85
Filtering ratio (%)	1.7	1.9	1.9




Figure 13. Gradient-based PIL analysis for NOAA AR 11121, illustrating that high-gradient structure concentrates along physically plausible polarity transitions in the super-resolved reconstruction.




Figure 14. Pixel-wise scatter of reconstructed LOS field from MagRes-Net versus the HMI reference over the clipped range $[-2000, 2000]$ G.




Figure 15. Frequency- and phase-domain fidelity of MagRes-Net.




Figure 16. Fourier phase error distribution.




Figure 17. Absolute reconstruction residual map showing the spatial distribution of $|\hat{B}_z - B_z|$.

provided for evaluation purposes. As can be seen from a comparison of the high resolution (HR) and super-resolution (SR) magnitude spectra in Fig. 15, there is good agreement in terms of their overall spectral distributions across all spatial frequencies; thus the super-resolution reconstruction has retained much of the large-scale organizational structure and many of the finer scale features present in the original HMI reference. Additionally, the frequency error heatmap in Fig. 15 demonstrates that the majority of the remaining errors occur primarily at lower frequency bands, whereas the higher frequency bands appear to have been preserved with very little error.




Figure 18. Noise robustness evaluation. PSNR and SSIM trends under increasing additive noise σ .

Table 7. Noise robustness under additive perturbation.

Noise σ	SSIM	PSNR (dB)
0	0.9777	44.31
5	0.7302	37.69
10	0.4326	32.72
15	0.2666	29.43
20	0.1761	27.00
25	0.1241	25.11
30	0.0919	23.55
35	0.0706	22.23
40	0.0562	21.11

We also evaluate the preservation of phase organization since phase information contained within the Fourier transform provides details on the spatial arrangement of the various components. Large phase errors could potentially result in structural deformities or artificial oscillations. In order to quantify this behaviour, phase maps and phase error images were generated and are presented in Fig. 15. These figures show that the phase organization across the vast majority of the frequency spectrum was preserved qualitatively when compared to the HMI reference. Statistical quantification of the preservation of phase organization is demonstrated by the phase difference histogram in Fig. 16, which shows that the majority of frequencies exhibit small phase errors, whereas a broader tail of larger phase errors correspond to localized areas of high gradient variation (e.g., near polarity reversal lines).

Finally, the spatial distribution of the errors is shown in Fig. 17, which presents an image of the absolute error $|\hat{B}_z - B_z|$ at each pixel. Residuals are not uniformly distributed, but instead are localized primarily in regions with strong magnetic structure and are much smaller in quiet background areas. Lastly we assess robustness to additive noise PSNR/SSIM in particular in Fig. 18 and Table 7. With increasing noise strength, both metrics degrade monotonically as to be expected from a stable, physically-consistent reconstruction framework, with PSNR quantitatively degrading from 44.31 dB @ $\sigma = 0$ to 21.11 dB @ $\sigma = 40$, similarly SSIM degrading from 0.9777, to 0.0562 over the same domain. That said, this performance behaviour neither yields abrupt degradation behaviour, meaning even strongly perturbative conditions yield meaningful reconstruction ability from the model. Spectral agreement, controlled phase deviation, spatially structured residuals and smooth relative drop-off from noise shows that SR is capable of preserving frequency space characteristics, and retaining physically agreeable stability against noisy conditions. A consolidated summary of the quantitative image-domain, perceptual, and spectral-domain evaluation metrics is provided in Table 8.

Table 8. Comprehensive quantitative evaluation of the proposed MagRes-Net reconstruction against the HMI reference magnetograms. The reported metrics quantify pixel-level accuracy, structural similarity, perceptual consistency, and spectral-domain fidelity.

Metric	Value
PSNR $\uparrow$ (dB)	44.31
SSIM $\uparrow$	0.97
RMSE $\downarrow$ (G)	2.76
LPIPS $\downarrow$	0.06
SAM $\downarrow$ (deg)	0.70
Frequency Correlation $\uparrow$	0.978
Phase Error Mean $\downarrow$ (rad)	1.4023
Phase Error Std. $\downarrow$ (rad)	0.9363

6.4. Ablation study

To quantify the respective roles of convolutional locality and Transformer-based global context modeling, we perform an ablation study using three model variants derived from the architecture shown in Fig. 3. The evaluated configurations include: (i) a ViT-only model without the CNN branch, (ii) a CNN-only model without the Transformer branch, and (iii) the full hybrid MagRes-Net architecture that integrates both branches through the feature fusion module. The quantitative comparison of these variants is summarized in Table 9.

The ViT-only variant relies exclusively on self-attention-based global representation learning. Transformer architectures are well suited for capturing long-range spatial dependencies and global contextual relationships (Dosovitskiy et al. 2021; Liu et al. 2021). Although the convolutional inductive bias was absent from this particular model, it was difficult for the model to achieve a good reconstruction of sharp polarity boundary locations and localized magnetic gradient magnitudes. Therefore, this specific model had the lowest PSNR and the largest RMSE of all models examined.

In comparison to the ViT-only version, the CNN-only version clearly outperformed it, which suggests that using localized convolutional processes has a strong advantage when attempting to recover relatively small scale, compact flux concentrations and fine-scale magnetic field topologies. The efficiency of CNNs at extracting spatially localized image features has been well documented within the literature related to image super-resolution methods (Dong et al. 2016; Lim et al. 2017; Zhang et al. 2018). Although the CNN-only model achieved an improvement in terms of reconstruction accuracy over the ViT-only model, the performance of the CNN-only model did not match that of the MagRes-Net model, which suggests that the use of solely localized processing is insufficient to provide a comprehensive description of larger scale configurations of magnetic fields in solar active regions.

The full MagRes-Net model allows the combination of the convolutional feature extraction process with the contextual modeling process provided by the Transformers; therefore, the network can provide both the reconstructed local magnetic gradients as well as preserve the overall large-scale structural properties of the solar active region. The results presented in Table 9 show that the PSNR is the greatest and the RMSE is the smallest for the hybrid model. This demonstrates that the inclusion of both local and global representations of features through the hybrid CNN-Transformer architecture illustrated in Figure 3 improves the quality of the reconstructed images. Furthermore, the ablation analysis illustrates that the hybrid CNN-Transformer model represents the best balance between the representation of the magnetogram super-resolution and the representation of the large-scale topology of the active region.

Table 9. Ablation analysis of the MagRes-Net architecture evaluating the individual contributions of convolutional local feature extraction and Transformer-based global context modeling. Performance is measured against the HMI reference magnetograms.

Model Variant	PSNR $\uparrow$ (dB)	RMSE $\downarrow$
ViT-only (global attention)	39.200	0.405
CNN-only (local gradients)	40.804	0.355
MagRes-Net (CNN + ViT)	44.31	0.276

6.5. Implications for flare-related magnetic diagnostics

Solar flare productivity is strongly linked to magnetic flux distribution, the complexity of polarity inversion lines, and the presence of strong magnetic gradients. Because these quantities are resolution sensitive, it is essential that super-resolution improves spatial detail while preserving physical consistency. The reconstructed maps preserve flux behaviour and enhance resolution-sensitive topology relevant to flare-feature computation (Schrijver 2007; Bobra and Couvidat 2015; Bobra et al. 2014; Florios et al. 2018; Nishizuka et al. 2018).

A central motivation for cross-instrument super-resolution is to ensure the consistency and continuity of flare-relevant magnetic diagnostics across solar cycles. Long-term investigations of solar variability and its coupling to the terrestrial atmosphere require magnetogram-derived quantities that are comparable over multi-decadal baselines. Many operational and research-grade flare predictors depend on diagnostics that are inherently resolution-sensitive, including unsigned magnetic flux Φ_{abs} , strong-field area coverage, and gradient-weighted polarity inversion line measures (Schrijver 2007; Bobra and Couvidat 2015; Bobra et al. 2014; Florios et al. 2018; Nishizuka et al. 2018). Resolution discrepancies between SOHO/MDI

and SDO/HMI therefore introduce systematic biases that propagate into long-term flare statistics and Sun–Earth interaction studies.

To connect super-resolution outputs to downstream diagnostic workflows without overstating forecasting performance, we evaluate whether MagRes-Net improves the diagnostic alignment of MDI-derived scalar quantities toward their HMI reference values. For any scalar diagnostic D computed over a region of interest, we define a scale-independent alignment score

$$\mathcal{A}(D) = 1 - \frac{|D_{\text{SR}} - D_{\text{HMI}}|}{|D_{\text{MDI}} - D_{\text{HMI}}| + \epsilon}, \quad (25)$$

where ϵ prevents division by zero. Values $\mathcal{A}(D) > 0$ indicate that the super-resolved estimate moves the diagnostic closer to the HMI reference relative to native MDI.

Beyond global flux budgets, polarity inversion line-driven measures are particularly resolution sensitive. A commonly used family of diagnostics weights PIL length by local magnetic-field gradient, reflecting the empirical association between strong-gradient PILs and flare productivity (Schrijver 2007; Georgoulis and Rust 2007). Using the extracted PIL set $\mathcal{L}$ obtained from Algorithm 1, we define a gradient-weighted PIL score

$$D_{\text{PIL}} = \sum_{(x,y) \in \mathcal{L}} \|\nabla B_z(x,y)\| \Delta\ell, \quad (26)$$

where $\Delta\ell$ denotes the pixel-length element along the PIL skeleton. When computed from native MDI, D_{PIL} is typically underestimated because spatial smoothing suppresses both gradient magnitude and fine-scale polarity reversals. Super-resolution that restores sharper polarity transitions is therefore expected to increase D_{PIL} toward HMI-consistent values.

Consistent with this expectation, the observed increase in filtered PIL length (Table 6), improved strong-field area coverage (Table 5), and enhanced magnetic stress proxy structure (Fig. 12) indicate that MagRes-Net restores resolution-dependent magnetic complexity while preserving polarity integrity. The coherent high-gradient PIL morphology shown in Fig. 13 further supports the conclusion that reconstructed maps recover physically meaningful topology rather than introducing artificial oscillatory artifacts.

From the perspective of atmospheric and space-climate research, improved cross-instrument homogeneity in magnetic diagnostics contributes to more reliable characterization of solar drivers of upper-atmospheric variability. Flare occurrence rates, magnetic complexity indices, and strong-gradient polarity inversion line statistics are frequently incorporated into empirical and physics-based models of ionospheric disturbance, thermospheric heating, and geomagnetic response. By reducing cross-cycle inconsistencies in these precursor quantities, MagRes-Net supports the construction of long-baseline, resolution-consistent solar forcing datasets. In this sense, the framework functions as a data-standardization layer that strengthens the integrity of Sun–atmosphere coupling analyses and facilitates more robust multi-decadal studies of solar–terrestrial interaction within the broader context of atmospheric sciences.

7. Conclusion and Future Work

The MagRes-Net, a novel hybrid super-resolution model based on Convolutional Neural Networks and Transformer, specifically designed to reduce the spatial resolution differences between SOHO/MDI and SDO/HMI magnetograms while preserving physically meaningful magnetic structure. The combined convolutional layers for local feature extraction with Transformer-based global contextual learning, MagRes-Net captures both small-scale magnetic gradients and larger-scale spatial relationships in solar active regions. A total of 5,232 co-temporal MDI–HMI magnetogram patch pairs were used to train the model. Image-level reconstruction metrics and physics-based diagnostic metrics were used to evaluate performance. Quantitatively, results on the held-out test set show that MagRes-Net significantly outperforms interpolation-based methods and CNN-only baselines. The proposed model achieves a PSNR of 44.31 dB and SSIM of 0.97, while reducing the reconstruction error to an RMSE of 2.76 G. In addition, perceptual and spectral evaluation metrics indicate strong structural fidelity, with LPIPS of 0.06, SAM of 0.70°, and a frequency-domain correlation of 0.978, demonstrating that the reconstructed magnetograms preserve both spatial structure and spectral characteristics of the HMI reference. Beyond improving pixel-level reconstruction accuracy, MagRes-Net provides notable gains in magnetic flux consistency, the morphology of compact strong-field active regions, and the structure of polarity inversion lines relative to interpolation-based and CNN-only models. A detailed analysis of NOAA Active Region AR 11121 further demonstrates enhanced recovery of compact flux concentrations and sharper magnetic gradients

that are substantially suppressed in the original MDI observations. Overall, these results demonstrate that enhancing existing magnetogram observations while reducing cross-instrument inconsistencies in current solar magnetic datasets. As such, MagRes-Net provides a pathway toward more homogeneous magnetogram data suitable for long-term studies of solar magnetic activity and flare-related diagnostics. Future work will include extending the evaluation of the proposed framework to a broader range of active regions and solar cycle phases. In addition, future efforts will focus on incorporating temporal modeling techniques to further improve the reliability of reconstructed data and facilitate longer-term solar activity analyses.

Acknowledgements The authors would like to acknowledge the National Institute of Technology Delhi for its support and resources used to complete this study. The authors also thank the science teams at SOHO/MDI and SDO/HMI for providing the publicly available magnetogram data essential to long-term solar magnetic field studies and space-weather analysis.

References

- Asensio Ramos, A., and Díaz Baso, C. J.: 2019, Bayesian deep learning for solar imaging. *Astron. Astrophys.*, **626**, A102.
- Bobra, M. G., Sun, X., Hoeksema, J. T., *et al.*: 2014, The Helioseismic and Magnetic Imager (HMI) SHARP data products. *Solar Phys.*, **289**, 3549–3578.
- Bobra, M. G., and Couvidat, S.: 2015, Solar flare prediction using SDO/HMI vector magnetic field data with a machine-learning algorithm. *Astrophys. J.*, **798**(2), 135.
- Díaz Baso, C. J., and Asensio Ramos, A.: 2018, A deep learning approach to solar image restoration. *Astron. Astrophys.*, **614**, A5.
- Dong, C., Loy, C. C., He, K., and Tang, X.: 2016, Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, **38**(2), 295–307.
- Dou, F., Xu, L., Ren, Z., Zhao, D., and Zhang, X.: 2022, Structure-preserving super-resolution for solar magnetograms. *Res. Astron. Astrophys.*, **22**(8), 085018.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., *et al.*: 2021, An image is worth 16×16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations (ICLR)*.
- Florios, K., Kontogiannis, I., Park, S.-H., *et al.*: 2018, Forecasting solar flares using magnetic field parameters and machine learning. *Solar Phys.*, **293**, 28.
- Georgoulis, M. K., and Rust, D. M.: 2007, Quantitative forecasting of major solar flares. *Astrophys. J. Lett.*, **661**(2), L109–L112.
- Hoeksema, J. T., Liu, Y., Hayashi, K., *et al.*: 2014, The Helioseismic and Magnetic Imager (HMI) vector magnetic field pipeline: Overview and performance. *Solar Phys.*, **289**, 3483–3530.
- Huang, A., Lu, G., Zhang, Y., Chen, S., Zhang, Y., Li, J., and Cummer, S. A.: 2018, Locating parent lightning strokes of sprites observed over a mesoscale convective system in Shandong Province, China. *Adv. Atmos. Sci.*, **35**, 1396–1414.
- Huang, H., Lu, G., An, A., Xu, D., Zhang, Y., Li, J., and Cummer, S. A.: 2025, Massive outbreak of red sprites in South Asia observed from the Tibetan Plateau. *Adv. Atmos. Sci.*, **42**, 1247–1260.
- Jarolim, R., Veronig, A. M., Pötzi, W., and Podladchikova, T.: 2025, Unpaired instrument-to-instrument translation for solar observations using diffusion models. *Nat. Commun.*, **16**, 3157.
- Jungbluth, A., *et al.*: 2019, Super-resolution of solar magnetograms for instrument bridging. arXiv:1911.01490.
- Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., and Timofte, R.: 2021, SwinIR: Image restoration using Swin Transformer. In: *Proceedings of the ICCV Workshops*.
- Lim, B., Son, S., Kim, H., Nah, S., and Lee, K. M.: 2017, Enhanced deep residual networks for single image super-resolution. In: *Proceedings of the CVPR Workshops*.
- Liu, Z., Lin, Y., Cao, Y., *et al.*: 2021, Swin Transformer: Hierarchical vision Transformer using shifted windows. In: *Proceedings of the ICCV*, 10012–10022.
- Mehra, P.: 1990, Association among geomagnetic activity, atmospheric electric field and selected meteorological parameters. *Adv. Atmos. Sci.*, **7**(2), 171–177.
- Munoz-Jaramillo, A., Jungbluth, A., Gitiaux, X., *et al.*: 2024, Superresolve and cross calibrate solar magnetograms. *Astrophys. J. Suppl. Ser.*, **271**(2), 46.
- Nishizuka, N., Sugiura, K., Kubo, Y., *et al.*: 2018, Solar flare prediction model with three machine-learning algorithms using active region magnetic parameters. *Astrophys. J.*, **858**, 113.
- Qie, X., Yuan, T., Zhang, Y., and Chen, C.: 2012, Progresses in atmospheric electricity research in China during 2006–2010. *Adv. Atmos. Sci.*, **29**, 993–1005.
- Qin, Y., *et al.*: 2024, Multi-scale enhancement for solar magnetograms. *Res. Astron. Astrophys.*, **24**(6), 065020.
- Rahman, S., *et al.*: 2020, Super-resolution strategies for solar magnetogram enhancement. *Astrophys. J. Lett.*, **897**(2), L32.
- Ramunno, F. P., *et al.*: 2025, Latent diffusion models for solar image super-resolution. arXiv:2503.24271.
- Scherrer, P. H., Bogart, R. S., Bush, R. I., *et al.*: 1995, The Solar Oscillations Investigation—Michelson Doppler Imager. *Solar Phys.*, **162**(1–2), 129–188.
- Scherrer, P. H., Schou, J., Bush, R. I., *et al.*: 2012, The Helioseismic and Magnetic Imager (HMI) investigation for the Solar Dynamics Observatory. *Solar Phys.*, **275**(1–2), 207–227.
- Schrijver, C. J.: 2007, A characteristic magnetic field pattern associated with all major solar flares and its use in flare forecasting. *Astrophys. J. Lett.*, **655**(2), L117–L120.
- Shi, W., Caballero, J., Huszár, F., *et al.*: 2016, Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: *Proceedings of the CVPR*, 1874–1883.
- Sikka, P., Selvam, A. M., and Murty, A. S. R.: 1988, Possible solar influence on atmospheric electric field. *Adv. Atmos. Sci.*, **5**(2), 217–228.
- Thompson, W. T.: 2006, Coordinate systems for solar image data. *Astron. Astrophys.*, **449**(2), 791–803.

-
- Wang, Q., Wu, B., Zhu, P., Li, P., Zuo, W., and Hu, Q.: 2020, ECA-Net: Efficient channel attention for deep convolutional neural networks. In: *Proceedings of the CVPR*, 11534–11542.
- Xu, X., He, S., and Wang, H.: 2020, Relationship between solar wind–magnetosphere energy and Eurasian winter cold events. *Adv. Atmos. Sci.*, **37**, 652–661.
- Xu, C., *et al.*: 2024, Attention-aided convolutional super-resolution for solar active-region magnetograms. *Solar Phys.*, **299**, 36.
- Zhang, Y., Li, Y., Li, K., Wang, L., Zhong, B., and Fu, Y.: 2018, Image super-resolution using very deep residual channel attention networks. In: *Proceedings of the ECCV*, 294–310.