
A Deep Learning Approach to Operational Flare Forecasting

Yasser Abdullah¹ Jason T. L. Wang²

Abstract

Solar flares are explosions on the Sun. They happen when energy stored in magnetic fields around solar active regions (ARs) is suddenly released. In this paper, we present a transformer-based framework, named SolarFlareNet, for predicting whether an AR would produce a γ -class flare within the next 24 to 72 hours. We consider three γ classes, namely the $\geq M5.0$ class, the $\geq M$ class and the $\geq C$ class, and build three transformers separately, each corresponding to a γ class. Each transformer is used to make predictions of its corresponding γ -class flares. The crux of our approach is to model data samples in an AR as time series and to use transformers to capture the temporal dynamics of the data samples. Each data sample consists of magnetic parameters taken from Space-weather HMI Active Region Patches (SHARP) and related data products. We survey flare events that occurred from May 2010 to December 2022 using the Geostationary Operational Environmental Satellite X-ray flare catalogs provided by the National Centers for Environmental Information (NCEI), and build a database of flares with identified ARs in the NCEI flare catalogs. This flare database is used to construct labels of the data samples suitable for machine learning. We further extend the deterministic approach to a calibration-based probabilistic forecasting method. The SolarFlareNet system is fully operational and is capable of making near real-time predictions of solar flares on the Web.

1. Introduction

Solar flares are sudden explosions of energy that occur on the Sun's surface. They often occur in solar active regions (ARs), caused by strong magnetic fields typically associated

with sunspot areas. Solar flares are categorized into five classes A, B, C, M, and X, with A-class flares having the lowest intensity and X-class flares having the highest intensity. Major flares are usually accompanied by coronal mass ejections and solar energetic particles (Qahwaji et al., 2008; Bobra & Ilonidis, 2016; Liu et al., 2017; Liu et al., 2020; Abdullah et al., 2022; Moreland et al., 2022). These eruptive events can have significant and harmful effects on or near Earth, damaging technologies, power grids, space stations, and human life (Daglis et al., 2004; Liu et al., 2019; Zhang et al., 2022; He et al., 2023). Therefore, providing accurate and early forecasts of solar flares is crucial for disaster risk management, risk mitigation, and preparedness.

Although a lot of effort has been devoted to flare prediction (Huang, X. et al., 2013; Panos & Kleint, 2020; Georgoulis et al., 2021; Tang et al., 2021), developing accurate, operational near-real-time flare forecasting systems remains a challenge. In the past, researchers designed statistical models for the prediction of flares based on the physical properties of active regions (Gallagher et al., 2002; Leka & Barnes, 2007; Mason & Hoeksema, 2010). With the availability of large amounts of flare-related data (Georgoulis et al., 2021), researchers started using machine learning methods for flare forecasting (Bobra & Couvidat, 2015; Liu et al., 2017; Abdullah et al., 2021a). More recently, deep learning, which is a subfield of machine learning, has emerged and showed promising results in predicting solar eruptions, including solar flares (Liu et al., 2022; Sun et al., 2022).

For example, Nishizuka et al. (Nishizuka et al., 2018) developed deep neural networks to forecast M- and C-class flares that would occur within 24 hours using data downloaded from the Solar Dynamics Observatory (SDO) (Pesnell, 2015) and the Geostationary Operational Environmental Satellite (GOES). Sun et al. (Sun et al., 2022) employed three-dimensional (3D) convolutional neural networks (CNNs) to forecast $\geq M$ -class and $\geq C$ -class flares using Space-weather HMI Active Region Patches (SHARP) (Bobra et al., 2014) magnetograms downloaded from the Joint Science Operations Center (JSOC) accessible at <http://jsoc.stanford.edu/>. Li et al. (Li et al., 2020) also adopted a CNN model to forecast $\geq M$ -class and $\geq C$ -class flares using SHARP magnetograms where the authors restructured the CNN layers in their neural network

¹Department of Computer Science, New Jersey Institute of Technology, Newark, NJ 07102, USA ²College of Computing, New Jersey Institute of Technology, Newark, NJ 07102, USA. Correspondence to: Jason T. L. Wang <wangj@njit.edu>.

with different filter sizes. Deng et al. (Deng et al., 2021) developed a hybrid CNN model to predict solar flares during the rising and declining phases of Solar Cycle 24.

Some researchers adopted SHARP magnetic parameters in time series for flare prediction. Static SHARP parameters quantitatively describe the properties of ARs, especially their ability to produce flares, at a given time. On the other hand, dynamic information, such as the magnetic helicity injection rate, sunspot motions, shear flows, and magnetic flux emergence/flux cancelation, is more important for flare forecasting. Using time series of SHARP parameters allows a model to capture the relationship between the evolution of magnetic fields of ARs and solar flares, hence achieving more accurate flare predictions (Tian, 2022; van Driel-Gesztelyi & Green, 2015). In an earlier study, Yu et al. (Yu et al., 2009) added the evolutionary information of ARs to a predictive model for the prediction of short-term solar flares. More recently, Chen et al. (Chen et al., 2019) designed a long short-term memory (LSTM) network to identify precursors of solar flare events using time series of SHARP parameters. LSTM is suitable for capturing the temporal dynamics of time series. Liu et al. (Liu et al., 2019) developed another LSTM network with a customized attention mechanism to direct the network to focus on important patterns in time series of SHARP parameters. Sun et al. (Sun et al., 2022) attempted to distinguish between ARs with strong flares ($\geq M$ -class flares) and ARs with no flare at all. The authors showed that combining LSTM and CNN can better solve the “strong versus quiet” flare prediction problem, with data from both Solar Cycle 23 and Cycle 24. All of the aforementioned studies provided valuable models and algorithms in the field. However, the existing methods focused on short-term forecasts (usually within 24 hours). Furthermore, the models were not used as operational systems.

In this paper, we propose a new deep learning approach to predicting solar flares using time series of SHARP parameters. Our approach employs a transformer-based framework, named SolarFlareNet, which predicts whether there would be a flare within 24 to 72 hours, where the flare could be a $\geq M5.0$ -, $\geq M$ -, or $\geq C$ -class flare. We further extend SolarFlareNet to produce probabilistic forecasts of flares and implement the probabilistic model into an operational, near real-time flare forecasting system. Experimental results demonstrate that SolarFlareNet generally performs better than, or is comparable to, related flare prediction methods.

2. Results

2.1. Deterministic Prediction Tasks

For any given active region (AR) and time point t , we predict whether there would be a γ -class flare within the next 24

hours (48 hours, 72 hours, respectively) of t where γ can be $\geq M5.0$, $\geq M$, or $\geq C$. A $\geq M5.0$ -class flare means the GOES X-ray flux value of the flare is above $5 \times 10^{-5} \text{Wm}^{-2}$. A $\geq M$ -class flare refers to an X- or M-class flare. A $\geq C$ -class flare refers to an X-class, M-class, or C-class flare. We focus on these three classes of flares due to their importance in space weather (Bobra & Couvidat, 2015; Jonas et al., 2018; Nishizuka et al., 2018; Liu et al., 2019). We developed three transformer models to tackle the three prediction tasks individually and separately. Notice that we did not consider γ to be $\geq X$ due to the lack of samples for X-class flares. Instead, we use $\geq M5.0$ as the most significant class, which contains sufficient samples.

2.2. Comparison with Previous Methods

We conducted a series of experiments to compare the proposed SolarFlareNet framework with closely related methods. All these methods perform binary classifications/predictions as defined above. We adopt several performance metrics. Formally, given an AR and a data sample x_t observed at time point t , we define x_t to be a true positive (TP) if the $\geq M5.0$ ($\geq M$, $\geq C$, respectively) model predicts that x_t is positive, i.e., the AR will produce a $\geq M5.0$ - ($\geq M$ -, $\geq C$ -, respectively) class flare within the next 24 hours of t , and x_t is indeed positive. We define x_t as a false positive (FP) if the $\geq M5.0$ ($\geq M$, $\geq C$, respectively) model predicts that x_t is positive while x_t is actually negative, i.e., the AR will not produce a $\geq M5.0$ - ($\geq M$ -, $\geq C$ -, respectively) class flare within the next 24 hours of t . We say x_t is a true negative (TN) if the $\geq M5.0$ ($\geq M$, $\geq C$, respectively) model predicts x_t to be negative and x_t is indeed negative; x_t is a false negative (FN) if the $\geq M5.0$ ($\geq M$, $\geq C$, respectively) model predicts x_t to be negative while x_t is actually positive. We also use TP (FP, TN, and FN, respectively) to represent the total number of true positives (false positives, true negatives, and false negatives, respectively). The TP, FP, TN, and FN for the 48-hour and 72-hour ahead predictions are defined similarly. The performance metrics are calculated as follows:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (1)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (2)$$

$$\text{Accuracy (ACC)} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{TN} + \text{FN}} \quad (3)$$

$$\text{Balanced ACC (BACC)} = \frac{\left(\frac{\text{TP}}{\text{TP} + \text{FN}} + \frac{\text{TN}}{\text{TN} + \text{FP}} \right)}{2} \quad (4)$$

$$\text{True Skill Statistics (TSS)} = \frac{\text{TP}}{\text{TP} + \text{FN}} - \frac{\text{FP}}{\text{FP} + \text{TN}} \quad (5)$$

Table 1 compares SolarFlareNet with related methods for 24-hour ahead flare predictions. The performance metric values of SolarFlareNet are mean values obtained from 10-fold cross-validation (Liu et al., 2019). The metric values of the highest performance models in the related studies are taken directly from those studies and are displayed in Table 1. The symbol ‘—’ means that a method does not produce the metric value for the corresponding prediction task. The best metric values are highlighted in boldface. TSS is the primary metric used in the literature on flare prediction. It can be seen from Table 1 that SolarFlareNet outperforms the state-of-the-art methods in terms of TSS except that Liu et al. (Liu et al., 2019) perform better than SolarFlareNet in predicting $\geq M5.0$ class flares.

Table 2 presents the mean performance metric values with standard deviations enclosed in parentheses for the 48- and 72-hour forecasts made by SolarFlareNet. None of the existing methods in Table 1 provides predictions in 48 or 72 hours in advance and, therefore, they are not listed in Table 2. Overall, SolarFlareNet performs well for the 48- and 72-hour forecasts. However, the metric values of the tool in Table 2 are lower than those in Table 1. This is understandable due to the longer range of predictions in Table 2.

2.3. Probabilistic Forecasting with Calibration

SolarFlareNet is essentially a probabilistic forecasting method, producing a probability between 0 and 1. The method compares the probability with a predetermined threshold, which is set to 0.5. Given an AR and a data sample x_t at time point t , if the predicted probability is greater than or equal to the threshold, then the AR will produce a flare within the next 24 (48, 72, respectively) hours of t (i.e., x_t belongs to the positive class); otherwise, the AR will not produce a flare within the next 24 (48, 72, respectively) hours of t (i.e., x_t belongs to the negative class). We can turn SolarFlareNet into a probabilistic forecasting method by directly outputting the predicted probability without comparing it with the threshold. Under this circumstance, the output is interpreted as a probabilistic estimate of how likely the AR will produce a flare within the next 24 (48, 72, respectively) hours of t . We employ a probability calibration technique with isotonic regression (Kruskal, 1964; Sager & Thisted, 1982) to adjust the predicted probability and avoid the mismatch between the distributions of the predicted and expected probabilistic values (Abduallah et al., 2022). We add a suffix “-C” to SolarFlareNet to denote the network without calibration.

To evaluate the performance of a probabilistic forecasting method, we use the Brier score (BS) and Brier skill score (BSS), defined as follows (Wilks, 2010; Liu et al., 2020;

Abduallah et al., 2022):

$$BS = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2 \quad (6)$$

$$BSS = 1 - \frac{BS}{\frac{1}{N} \sum_{i=1}^N (y_i - \bar{y})^2} \quad (7)$$

where N is the number of data samples in a test set; y_i denotes the observed probability and \hat{y}_i denotes the predicted probability of the i th test data sample, respectively; $\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i$ denotes the mean of all the observed probabilities. BS values range from 0 to 1, with 0 being a perfect score. BSS values range from $-\infty$ to 1, with 1 being a perfect score.

Table 3 compares SolarFlareNet, used as a probabilistic forecasting method, with a closely related method (Liu et al., 2019). The BS and BSS values in the table are mean values (with standard deviations enclosed in parentheses) obtained from 10-fold cross-validation. The metric values for the existing method are taken directly from the related work (Liu et al., 2019). The best BS and BSS values are highlighted in bold. Notice that the existing method did not make 48-hour or 72-hour forecasts in advance. Table 3 shows that there is no definitive conclusion regarding the relative performance of SolarFlareNet and the existing method. The existing method is better in terms of BS, while SolarFlareNet is better in terms of BSS. On the other hand, the calibrated version of a model is better than the model without calibration. Notice also that the results of the 48-hour and 72-hour forecasts are worse than those of the 24-hour forecasts. This is understandable since the longer the prediction window, the worse the performance a model achieves due to data deviation over time.

2.4. The SolarFlareNet System

We have implemented the probabilistic forecasting method described above into an operational, near real-time flare forecasting system. To access the system, visit the SolarDB website at <https://nature.njit.edu/solardb/index.html>. On the website, select and click the menu entry “Tools” and then select and click “Flare Forecasting System.” Figure 1 shows the graphical user interface (GUI) of the system. It displays a probabilistic estimate of how likely an AR will produce a flare within the next 24, 48, and 72 hours of the time point at which the system is updated each day. No prediction is made for an AR marked with a special character *, #, or ~ where

- * means the AR is near the limb,
- # means the AR is spotless with the number of spots being zero,

Metric	Method	$\geq M5.0$ class	$\geq M$ class	$\geq C$ class
Recall	Huang et al. (Huang et al., 2018)	—	—	—
	Li et al. (Li et al., 2020)	—	0.817	0.889
	Liu et al. (Liu et al., 2019)	0.960	0.885	0.773
	Sun et al. (Sun et al., 2022)	—	0.925	0.862
	Wang et al. (Wang et al., 2020)	—	0.730	0.621
	This work	0.853	0.842	0.891
Precision	Huang et al. (Huang et al., 2018)	—	—	—
	Li et al. (Li et al., 2020)	—	0.889	0.906
	Liu et al. (Liu et al., 2019)	0.048	0.222	0.541
	Sun et al. (Sun et al., 2022)	—	0.595	0.878
	Wang et al. (Wang et al., 2020)	—	0.282	0.541
	This work	0.977	0.848	0.949
ACC	Huang et al. (Huang et al., 2018)	—	—	—
	Li et al. (Li et al., 2020)	—	0.891	0.861
	Liu et al. (Liu et al., 2019)	0.921	0.907	0.826
	Sun et al. (Sun et al., 2022)	—	0.904	0.879
	Wang et al. (Wang et al., 2020)	—	0.945	0.883
	This work	0.964	0.928	0.915
BACC	Huang et al. (Huang et al., 2018)	—	—	—
	Li et al. (Li et al., 2020)	—	—	—
	Liu et al. (Liu et al., 2019)	0.940	0.896	0.806
	Sun et al. (Sun et al., 2022)	—	—	—
	Wang et al. (Wang et al., 2020)	—	—	—
	This work	0.926	0.919	0.917
TSS	Huang et al. (Huang et al., 2018)	—	0.662	0.487
	Li et al. (Li et al., 2020)	—	0.749	0.679
	Liu et al. (Liu et al., 2019)	0.881	0.792	0.612
	Sun et al. (Sun et al., 2022)	—	0.826	0.756
	Wang et al. (Wang et al., 2020)	—	0.681	0.553
	This work	0.818	0.839	0.835

Table 1. Performance comparison between SolarFlareNet and related methods for 24-hour ahead flare predictions.

Metric	Hour	$\geq M5.0$ class	$\geq M$ class	$\geq C$ class
Recall	48	0.739 (0.108)	0.735 (0.089)	0.722 (0.089)
	72	0.717 (0.100)	0.708 (0.078)	0.702 (0.089)
Precision	48	0.890 (0.210)	0.823 (0.092)	0.812 (0.072)
	72	0.872 (0.045)	0.812 (0.089)	0.809 (0.051)
ACC	48	0.923 (0.003)	0.907 (0.007)	0.896 (0.047)
	72	0.906 (0.002)	0.883 (0.005)	0.863 (0.040)
BACC	48	0.864 (0.054)	0.857 (0.045)	0.848 (0.040)
	72	0.856 (0.039)	0.843 (0.048)	0.834 (0.029)
TSS	48	0.736 (0.112)	0.728 (0.090)	0.719 (0.079)
	72	0.729 (0.108)	0.714 (0.095)	0.709 (0.058)

Table 2. Performance metric values of SolarFlareNet for 48- and 72-hour ahead flare predictions.

A Deep Learning Approach to Operational Flare Forecasting

Hour	Metric	Method	$\geq M5.0$ class	$\geq M$ class	$\geq C$ class
24	BS	Liu et al. (Liu et al., 2019)	0.090 (0.011)	0.090 (0.009)	0.133 (0.007)
		SolarFlareNet	0.226 (0.024)	0.244 (0.013)	0.285 (0.034)
		SolarFlareNet-C	0.263 (0.024)	0.281 (0.050)	0.313 (0.033)
	BSS	Liu et al. (Liu et al., 2019)	-21.576 (2.956)	-2.241 (0.319)	0.152 (0.047)
		SolarFlareNet	0.584 (0.022)	0.521 (0.042)	0.409 (0.062)
		SolarFlareNet-C	0.504 (0.026)	0.491 (0.031)	0.349 (0.055)
48	BS	Liu et al. (Liu et al., 2019)	—	—	—
		SolarFlareNet	0.272 (0.091)	0.312 (0.101)	0.361 (0.091)
		SolarFlareNet-C	0.315 (0.049)	0.336 (0.033)	0.378 (0.104)
	BSS	Liu et al. (Liu et al., 2019)	—	—	—
		SolarFlareNet	0.569 (0.045)	0.524 (0.021)	0.502 (0.033)
		SolarFlareNet-C	0.457 (0.062)	0.424 (0.091)	0.411 (0.056)
72	BS	Liu et al. (Liu et al., 2019)	—	—	—
		SolarFlareNet	0.313 (0.062)	0.327 (0.063)	0.344 (0.049)
		SolarFlareNet-C	0.329 (0.094)	0.369 (0.088)	0.376 (0.102)
	BSS	Liu et al. (Liu et al., 2019)	—	—	—
		SolarFlareNet	0.549 (0.067)	0.524 (0.089)	0.501 (0.093)
		SolarFlareNet-C	0.514 (0.077)	0.469 (0.095)	0.447 (0.059)

Table 3. Performance comparison between SolarFlareNet and an existing method for probabilistic flare predictions (24 to 72 hours in advance).

- \sim means no SHARP data is available for the AR.

The system provides daily predictions based on the data obtained from the previous day. When the user clicks the link to the previous day, the user is led to the SolarMonitor site that is accessible at <https://www.solarmonitor.org/index.php> where detailed AR information for that day is available. The system also provides previous forecasting results since the operational system came online. We compare the previous forecasting results with the true flare events in the GOES X-ray flare catalogs provided by NCEI. The SolarFlareNet system achieves 89% (76%, 71%, respectively) accuracy for 24-hour (48-hour, 72-hour, respectively) ahead predictions.

3. Discussion and Conclusion

In this paper, we present a novel transformer-based framework to predict whether a solar active region (AR) would produce a γ -class flare within the next 24 to 72 hours where γ is $\geq M5.0$, $\geq M$, or $\geq C$. We use three transform models to handle the three classes of flares individually and separately. All three transformer models perform binary predictions. We collect ARs with flares that occurred between 2010 and 2022 from the GOES X-ray flare catalogs provided by the National Centers for Environmental Information (NCEI). In addition, we downloaded SHARP magnetic parameters from the Joint Science Operations Center (JSOC). Each data

sample contains SHARP parameters suitable for machine learning. We conducted experiments using 10-fold cross-validation (Liu et al., 2019). Based on the experiments, our transformer-based framework generally performs better than closely related methods in terms of TSS (true skill statistics), as shown in Table 1. We further extend our framework to produce probabilistic forecasts of flares and implement the framework into an operational, near real-time flare forecasting system accessible on the Web. The probabilistic framework is comparable to a closely related method (Liu et al., 2019) in terms of BS (Brier score) and BSS (Brier skill score) when making 24-hour forecasts, as shown in Table 3, although the existing method did not make 48- or 72-hour forecasts. Thus, we conclude that SolarFlareNet is a feasible tool for producing flare forecasts within 24 to 72 hours.

4. Methods

4.1. Data Collection

In this study we used SHARP magnetic parameters (Bobra et al., 2014; Bobra & Ilonidis, 2016; Liu et al., 2019) downloaded from the Joint Science Operations Center (JSOC) accessible at <http://jsoc.stanford.edu/>. Specifically, we collect data samples, composed of SHARP parameters, at a cadence of 12 minutes where the data samples are retrieved from the `hmi.sharp_cea_720s` data series on

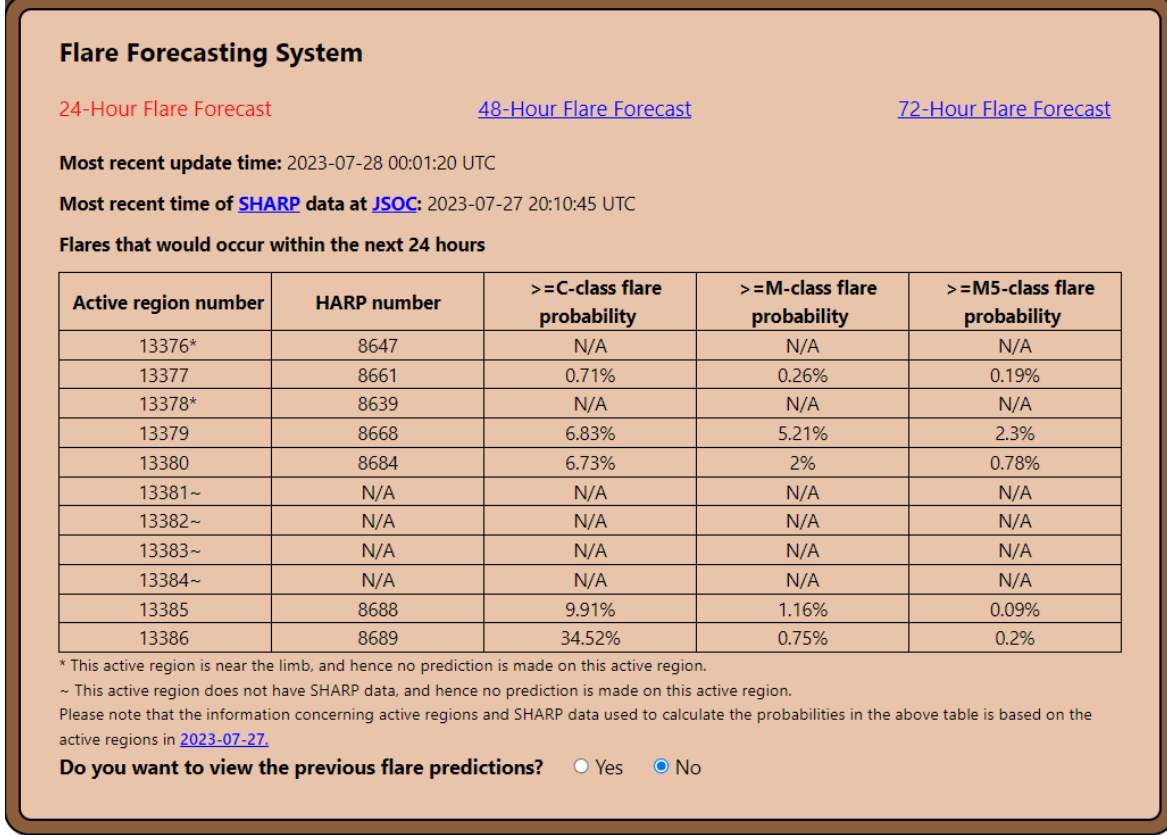


Figure 1. The graphical user interface of the SolarFlareNet system.

the JSOC website using the Python package SunPy (The SunPy Community et al., 2015). We selected nine SHARP magnetic parameters as suggested in the literature (Bobra & Couvidat, 2015; Bobra & Ilonidis, 2016; Liu et al., 2017; 2019; Liu et al., 2020). These nine parameters include the total unsigned current helicity (TOTUSJH), total unsigned vertical current (TOTUSJZ), total unsigned flux (USFLUX), mean characteristic twist parameter (MEANALP), sum of flux near polarity inversion line (R_VALUE), total photospheric magnetic free energy density (TOTPOT), sum of the modulus of the net current per polarity (SAVNCPP), area of strong field pixels in the active region (AREA_ACR), and absolute value of the net current helicity (ABS NJZH). Table 4 presents an overview of the nine parameters. The SHARP parameters' values are in different scales and units; therefore, we normalize each parameter's values using the min-max normalization method (Liu et al., 2020; Abdullallah et al., 2022). Formally, let p_i^k be the original value of the i th parameter of the k th data sample. Let q_i^k be the normalized value of the i th parameter of the k th data sample. Let min_i be the minimum value of the i th parameter. Let max_i be the maximum value of the i th parameter. Then

$$q_i^k = \frac{p_i^k - min_i}{max_i - min_i} \quad (8)$$

We collected A-, B-, C-, M- and X-class flares that occurred between May 2010 and December 2022, and their associated active regions (ARs) from the GOES X-ray flare catalogs provided by the National Centers for Environmental Information (NCEI). Flares without identified ARs were excluded. This process yielded a database of 8 A-class flares, 6,571 B-class flares, 8,973 C-class flares, 895 M-class flares, and 58 X-class flares. Also, we collected 10 nonflaring ARs (Hazra et al., 2020). We collected data samples that were 24 (48, 72, respectively) hours before a flare. Furthermore, we collected data samples that were 24 (48, 72, respectively) hours after the start time of each nonflaring AR. The data was then cleaned as follows (Bobra & Ilonidis, 2016; Liu et al., 2019; Abdullallah et al., 2022).

We discard ARs that are outside $\pm 70^\circ$ of the central meridian. These ARs are near the limb and have projection effects that render the calculation of the ARs' SHARP parameters incorrect. In addition, we discard a data sample if (i) its corresponding flare record has an absolute value of the radial velocity of SDO greater than 3500 m s^{-1} , (ii) the HMI data have low quality (Hoeksema et al., 2014), or (iii) the data sample has missing values or incomplete SHARP parameters. Thus, we exclude low-quality data samples and keep qualified data samples of high quality in our study.

Keyword	Description	Formula
TOTUSJH	Total unsigned current helicity	$H_{ctotal} \propto \sum B_z \cdot J_z $
TOTUSJZ	Total unsigned vertical current	$J_{ztotal} = \sum J_z dA$
USFLUX	Total unsigned flux	$\Phi = \sum B_z dA$
MEANALP	Mean characteristic twist parameter, α	$\alpha_{total} \propto \frac{\sum J_z B_z}{\sum B_z^2}$
R.VALUE	Sum of flux near polarity inversion line	$\Phi = \sum B_{LOS} dA$ within R mask
TOTPOT	Total photospheric magnetic free energy density	$\rho_{tot} \propto \sum (\mathbf{B}^{Obs} - \mathbf{B}^{Pot})^2 dA$
SAVNCPP	Sum of the modulus of the net current per polarity	$J_{zsum} \propto \sum^{B_z^+} J_z dA + \sum^{B_z^-} J_z dA $
AREA_ACR	Area of strong field pixels in the active region	Area = \sum Pixels
ABSNJZH	Absolute value of the net current helicity	$H_{cabs} \propto \sum B_z \cdot J_z $

Table 4. Overview of the nine SHARP parameters used in our study.

4.2. Data Labeling

Data labeling is crucial in machine learning. To predict $\geq C$ -class flares, suppose that a C-, M-, or X-class flare occurs at time point t on an AR (more precisely, the start time of the flare is t). Data samples between t and $t - 24$ hours (48, 72 hours, respectively) in the AR are labeled positive. If the flare occurs at time point t is an A-class or B-class flare, the data samples between t and $t - 24$ hours (48, 72 hours, respectively) in the AR are labeled negative. Figure 2 illustrates the labeling scheme to predict whether a $\geq C$ -class flare would occur within 24 hours. In predicting $\geq M$ -class flares, we use $\geq M$ -class flares to label positive data samples; use $\leq C$ -class flares to label negative data samples. In predicting $\geq M5.0$ -class flares, we use $\geq M5.0$ -class flares to label positive data samples; use $\leq C$ -class flares as well as M1.0- through M4.0-class flares to label negative data samples. If there are recurring flares whose corresponding data samples overlap, we give priority to the largest flare and label the overlapped data samples based on the largest flare. In all three prediction tasks, the data samples in the nonflaring ARs are labeled negative.

Table 5 shows the total numbers of positive and negative data samples in each class for 24-, 48-, and 72-hour ahead flare predictions. The numbers in the table are lower than expected. This is because we discarded/removed many low-quality data samples as described above. If a gap occurs in the middle of a time series due to removal, we use a zero-padding strategy (Liu et al., 2019; Abduallah et al., 2022) to create a synthetic data sample to fill the gap. The synthetic data sample has zero values for all nine SHARP parameters. The synthetic data sample is added after normalization of the values of the SHARP parameters, and therefore the synthetic data sample does not affect the normalization procedure.

For each prediction task, we divide the corresponding data samples into 10 equal sized distinct partitions/folds that are used to perform 10-fold cross-validation experiments. In the run i , where $1 \leq i \leq 10$, we use fold i as the test set and the union of the other nine folds as the training set. The

data samples of the same AR are placed in the training set or the test set, but not both. This scheme ensures that a model is trained with data different from the test data and makes predictions on the test data that it has never seen during training. There are 10 folds and, consequently, 10 runs. The means and standard deviations of the performance metrics' values over the 10 runs are calculated and recorded.

4.3. Data Augmentation

The data sets used in this study to predict flares of the $\geq M$ - and $\geq M5.0$ -class are imbalanced as shown in Table 5 where negative data samples are much more than positive data samples. Imbalanced data pose a challenge in model training and often result in poor model performance. One may use data augmentation to combat the imbalanced data. Data augmentation is an important technique that enriches training data and increases the generalization of the model (Deng et al., 2021). Here, we adopt the Gaussian white noise (GWN) data augmentation scheme because it has shown a significant improvement in model performance (Um et al., 2017; Li et al., 2020). GWN assumes that any two values are statistically independent, regardless of how close they are in time. The stationary random values of GWN are generated using the zero mean and 5% of the standard deviation. During training, the data augmentation is applied to the minority (positive) class, leaving the majority (negative) class as is. During testing, the data are left without any augmentation so that the model predicts only on the actual test data to avoid any misleading performance assessment.

4.4. The SolarFlareNet Architecture

Figure 3 presents the architecture of SolarFlareNet. It is a transformer-based framework that combines a one-dimensional convolutional neural network (Conv1D), long short-term memory (LSTM), transformer encoder blocks (TEBs), and additional layers that include batch normalization (BN) layers, dropout layers, and dense layers. The first layer is the input layer, which takes as input a time series of m consecutive data samples $x_{t-m+1}, x_{t-m+2} \dots x_{t-1}$,




Figure 2. Illustration of positive and negative data samples used in predicting $\geq C$ -class flares. In the left panel, the red vertical line indicates the start time of a $\geq C$ -class flare. The data samples collected in the 24 hours prior to the red vertical line are labeled positive (in green color). In the right panel, the red vertical line indicates the start time of an A-class or B-class flare. The data samples collected in the 24 hours prior to the red vertical line are labeled negative (in yellow color).

Hour	Data samples	$\geq M5.0$ class	$\geq M$ class	$\geq C$ class
24	Positive	2,125	13,989	244,968
	Negative	461,060	449,196	218,517
48	Positive	2,255	16,709	316,149
	Negative	615,708	602,154	304,714
72	Positive	2,375	18,505	356,219
	Negative	704,997	689,567	350,953

Table 5. Total numbers of positive and negative data samples in each class for 24-, 48-, and 72-hour ahead flare predictions.

x_t where x_t is the data sample at time point t (Abduallah et al., 2022). (In the study presented here, m is set to 10.) The input layer is followed by a BN layer. BN is an additional mechanism to stabilize SolarFlareNet, make it faster, and help to avoid overfitting during training (Zerveas et al., 2021). We applied BN after the input layer, the LSTM layer, and within the TEBs to make sure that SolarFlareNet is stable throughout the training process. The BN layer is followed by the Conv1D layer because time series generally have a strong 1D time locality that can be extracted by the Conv1D layers (Kravchik & Shabtai, 2018). Then, the LSTM layer is used, which is equipped with regularization to also avoid overfitting. LSTM is suitable for handling time series data to capture the temporal correlation and dependency in the data. Adding an LSTM layer after a Conv1D layer has shown significant improvement in time series prediction (Abduallah et al., 2021b; 2022b;a). The LSTM layer passes the learned features and patterns to a BN layer to stabilize the network before the data go to the TEBs.

We use transformer encoders without decoders because we process time series here, rather than performing natural language processing where the decoders are required to decode the words for sentence translation. The number of TEBs is set to 4. This number has a significant effect on the overall performance of the model (Vaswani et al., 2017). When we use less than 4 TEBs, the model is not able to learn useful patterns and is under-fitted. When we

use more than 4 TEBs, the large number of TEBs causes overhead on the encoder processing while the model tends to do excessive overfitting and lean toward the majority class (i.e., negative class) in the data, ignoring the minority class (i.e., positive class) entirely. Each TEB is configured with a dropout layer, multi-head attention (MHA) layer, a BN layer, a Conv1D layer, and an LSTM layer. The MHA layer is the most important layer in the encoder because it provides the transformation on the sequence values to obtain the different metrics. The MHA layer is configured with 4 heads and each attention head is also set to 4. The dropout layer is mainly used to overcome the overfitting caused by the imbalanced data. It drops a percentage of the neurons from the architecture, which causes the internal architecture of the model to change, allowing for better performance and stability. Finally, the softmax function is used as the final activation function, which produces a probabilistic estimate of how likely a flare will occur within the next 24 (48, 72, respectively) hours of t .

4.5. Ablation Study

We performed ablation tests to assess each component of SolarFlareNet. We consider four variants of SolarFlareNet, denoted SolarFlareNet-Conv, SolarFlareNet-L, SolarFlareNet-ConvL, and SolarFlareNet-T, respectively. Here, SolarFlareNet-Conv (SolarFlareNet-L, SolarFlareNet-ConvL, SolarFlareNet-T, respectively) represents the subnet

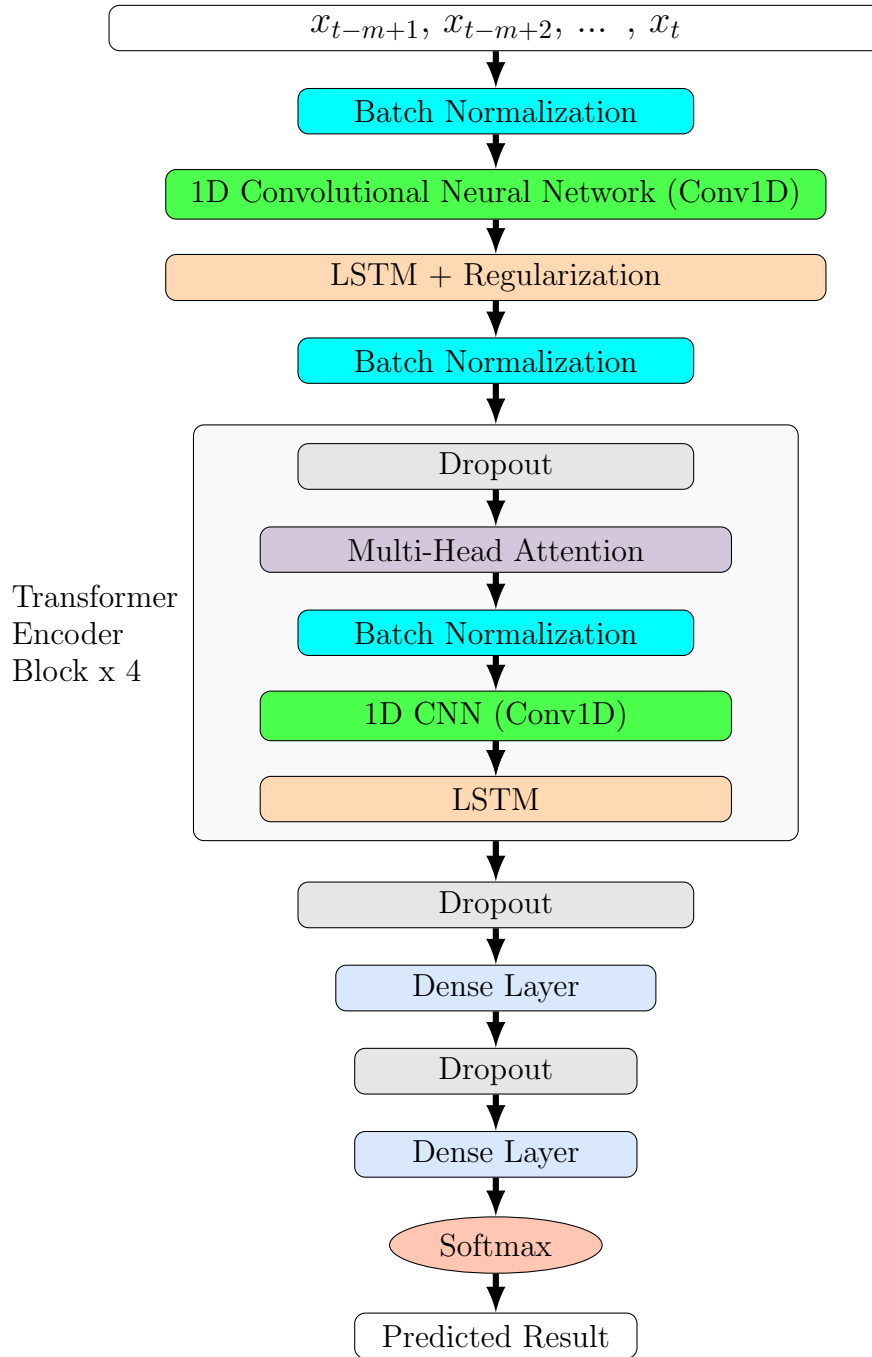


Figure 3. Architecture of SolarFlareNet.

of SolarFlareNet in which the Conv1D layer (LSTM layer, Conv1D and LSTM layers, transformer network with the 4 TEBs, respectively) is removed while keeping the remaining components of the SolarFlareNet framework. Table 6 compares the TSS values of the five models for the 24-, 48-, and 72-hour ahead flare prediction. It can be seen from Table 6 that the full model, SolarFlareNet, outperforms the four subnets in terms of the TSS metric. This happens be-

cause the SolarFlareNet-Conv model captures the temporal correlation of the test data, but does not learn additional characteristics of the data to build a stronger relationship between the test data. SolarFlareNet-L captures the properties of the test data, but lacks knowledge of the temporal correlation patterns in the data to deeply analyze the sequential information in the test data. It can also be seen from Table 6 that the SolarFlareNet-ConvL model is not as

A Deep Learning Approach to Operational Flare Forecasting

Hour	Method	$\geq M5.0$ class	$\geq M$ class	$\geq C$ class
24	SolarFlareNet	0.818 (0.021)	0.839 (0.030)	0.835 (0.048)
	SolarFlareNet-Conv	0.780 (0.036)	0.759 (0.052)	0.822 (0.023)
	SolarFlareNet-L	0.779 (0.022)	0.737 (0.041)	0.713 (0.039)
	SolarFlareNet-ConvL	0.742 (0.029)	0.719 (0.041)	0.728 (0.037)
	SolarFlareNet-T	0.716 (0.101)	0.704 (0.093)	0.712 (0.078)
48	SolarFlareNet	0.736 (0.112)	0.728 (0.090)	0.719 (0.079)
	SolarFlareNet-Conv	0.729 (0.049)	0.715 (0.055)	0.695 (0.035)
	SolarFlareNet-L	0.694 (0.066)	0.689 (0.012)	0.675 (0.021)
	SolarFlareNet-ConvL	0.681 (0.063)	0.676 (0.054)	0.673 (0.048)
	SolarFlareNet-T	0.662 (0.061)	0.647 (0.032)	0.641 (0.033)
72	SolarFlareNet	0.729 (0.108)	0.714 (0.095)	0.709 (0.058)
	SolarFlareNet-Conv	0.703 (0.042)	0.696 (0.011)	0.658 (0.023)
	SolarFlareNet-L	0.688 (0.046)	0.666 (0.039)	0.658 (0.016)
	SolarFlareNet-ConvL	0.665 (0.026)	0.643 (0.031)	0.632 (0.030)
	SolarFlareNet-T	0.635 (0.028)	0.624 (0.046)	0.619 (0.033)

Table 6. TSS values of the five methods considered in the ablation study.

good as the full model, indicating that the transformer network alone is not enough to produce the best results. Lastly, SolarFlareNet-T has the least performance among the four subnets, demonstrating the importance of the transformer network. In conclusion, our ablation study indicates that the performance of the proposed SolarFlareNet framework is not dominated by any single component. In fact, all components have made contributions to the overall performance of the proposed framework.

References

- Abduallah, Y., Wang, J. T. L., Nie, Y., Liu, C., and Wang, H. DeepSun: Machine-learning-as-a-service for solar flare prediction. *Research in Astronomy and Astrophysics*, 21(7):160, Aug 2021a. URL <https://doi.org/10.1088/1674-4527/21/7/160>.
- Abduallah, Y., Wang, J. T. L., Shen, Y., Alobaid, K. A., Criscuoli, S., and Wang, H. Reconstruction of total solar irradiance by deep learning. In *Proceedings of the 34th International Florida Artificial Intelligence Research Society Conference*, 2021b. URL <https://doi.org/10.32473/flairs.v34i1.128356>.
- Abduallah, Y., Jordanova, V. K., Liu, H., Li, Q., Wang, J. T. L., and Wang, H. Predicting solar energetic particles using SDO/HMI vector magnetic data products and a bidirectional LSTM network. *The Astrophysical Journal Supplement Series*, 260(1):16, May 2022.
- Abduallah, Y., Wang, J. T. L., Bose, P., Zhang, G., Gerges, F., and Wang, H. Forecasting the disturbance storm time index with Bayesian deep learning. In *Proceedings of the 35th International Florida Artificial Intelligence Research Society Conference*, 2022a. URL <https://doi.org/10.32473/flairs.v35i.130564>.
- Abduallah, Y., Wang, J. T. L., Xu, C., and Wang, H. A transformer-based framework for geomagnetic activity prediction. In *Proceedings of the 26th International Symposium on Methodologies for Intelligent Systems*, Oct. 2022b. ISBN 978-3-031-16563-4. URL https://doi.org/10.1007/978-3-031-16564-1_31.
- Bobra, M. G. and Couvidat, S. Solar flare prediction using SDO/HMI vector magnetic field data with a machine-learning algorithm. *The Astrophysical Journal*, 798:135, Jan 2015.
- Bobra, M. G. and Ilonidis, S. Predicting coronal mass ejections using machine learning methods. *The Astrophysical Journal*, 821(2):127, April 2016.
- Bobra, M. G., Sun, X., Hoeksema, J. T., Turmon, M., Liu, Y., Hayashi, K., Barnes, G., and Leka, K. D. The Helioseismic and Magnetic Imager (HMI) vector magnetic field pipeline: SHARPs - Space-weather HMI Active Region Patches. *Solar Physics*, 289:3549–3578, September 2014.
- Chen, Y., Manchester, W. B., Hero, A. O., Toth, G., DuFumier, B., Zhou, T., Wang, X., Zhu, H., Sun, Z., and Gombosi, T. I. Identifying solar flare precursors using time series of SDO/HMI images and SHARP parameters. *Space Weather*, 17(10):1404–1426, 2019.
- Daglis, I., Baker, D., Kappenman, J., Panasyuk, M., and Daly, E. Effects of space weather on technology infrastructure. *Space Weather*, 2:S02004, 2004.
- Deng, Y., Lu, L., Aponte, L., Angelidi, A. M., Novak, V., Karniadakis, G. E., and Mantzoros, C. S. Deep transfer

- learning and data augmentation improve glucose levels prediction in type 2 diabetes patients. *Nature Portfolio Journal Digital Medicine*, 4:1013345, July 2021.
- Deng, Z., Wang, F., Deng, H., Tan, L., Deng, L., and Feng, S. Fine-grained solar flare forecasting based on the hybrid convolutional neural networks. *The Astrophysical Journal*, 922(2):232, December 2021.
- Gallagher, P. T., Moon, Y. J., and Wang, H. Active-region monitoring and flare forecasting I. Data processing and first results. *Solar Physics*, 209(1):171–183, September 2002.
- Georgoulis, M. K., Bloomfield, D. S., Piana, M., Massone, A. M., Soldati, M., Gallagher, P. T., Pariat, E., Vilmer, N., Buchlin, E., Baudin, F., Csillaghy, A., Sathiapal, H., Jackson, D. R., Alingery, P., Benvenuto, F., Campi, C., Florios, K., Gontikakis, C., Guennou, C., Guerra, J. A., Kontogiannis, I., Latorre, V., Murray, S. A., Park, S.-H., von Stachelski, S., Torbica, A., Vischi, D., and Worsfold, M. The flare likelihood and region eruption forecasting (FLARECAST) project: flare forecasting in the big data & machine learning era. *J. Space Weather Space Climate*, 11:39, 2021. URL <https://doi.org/10.1051/swsc/2021023>.
- Hazra, S., Sardar, G., and Chowdhury, P. Distinguishing between flaring and nonflaring active regions. *Astronomy & Astrophysics*, 639:A44, July 2020.
- He, X.-r., Zhong, Q.-z., Cui, Y.-m., Liu, S.-q., Shi, Y.-r., Yan, X.-h., and Wang, Z.-s.-y. Solar flare short-term forecast model based on long and short-term memory neural network. *Chinese Astronomy and Astrophysics*, 47(1):108–126, January 2023.
- Hoeksema, J. T., Liu, Y., Hayashi, K., Sun, X., Schou, J., Couvidat, S., Norton, A., Bobra, M., Centeno, R., Leka, K. D., Barnes, G., and Turmon, M. The Helioseismic and Magnetic Imager (HMI) vector magnetic field pipeline: Overview and performance. *Solar Physics*, 289(9):3483–3530, September 2014.
- Huang, X., Wang, H., Xu, L., Liu, J., Li, R., and Dai, X. Deep learning based solar flare forecasting model. I. Results for line-of-sight magnetograms. *The Astrophysical Journal*, 856(1):7, March 2018.
- Huang, X., Zhang, L., Wang, H., and Li, L. Improving the performance of solar flare prediction using active longitudes information. *Astronomy & Astrophysics*, 549:A127, 2013. URL <https://doi.org/10.1051/0004-6361/201219742>.
- Jonas, E., Bobra, M., Shankar, V., Todd Hoeksema, J., and Recht, B. Flare prediction using photospheric and coronal image data. *Solar Physics*, 293(3):48, March 2018.
- Kravchik, M. and Shabtai, A. Detecting cyber attacks in industrial control systems using convolutional neural networks. In *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy*, 2018. ISBN 9781450359924. URL <https://doi.org/10.1145/3264888.3264896>.
- Kruskal, J. B. Nonmetric multidimensional scaling: A numerical method. *Psychometrika*, 29(2):115–129, 1964.
- Leka, K. D. and Barnes, G. Photospheric magnetic field properties of flaring versus flare-quiet active regions. IV. A statistically significant sample. *The Astrophysical Journal*, 656(2):1173–1186, February 2007.
- Li, K., Daniels, J., Liu, C., Herrero, P., and Georgiou, P. Convolutional recurrent neural networks for glucose prediction. *IEEE Journal of Biomedical and Health Informatics*, 24(2):603–613, 2020.
- Li, X., Zheng, Y., Wang, X., and Wang, L. Predicting solar flares using a novel deep convolutional neural network. *The Astrophysical Journal*, 891(1):10, March 2020.
- Liu, C., Deng, N., Wang, J. T. L., and Wang, H. Predicting solar flares using SDO/HMI vector magnetic data products and the random forest algorithm. *The Astrophysical Journal*, 843:104, 2017.
- Liu, H., Liu, C., Wang, J. T. L., and Wang, H. Predicting solar flares using a long short-term memory network. *The Astrophysical Journal*, 877(2):121, jun 2019.
- Liu, H., Liu, C., Wang, J. T. L., and Wang, H. Predicting coronal mass ejections using SDO/HMI vector magnetic data products and recurrent neural networks. *The Astrophysical Journal*, 890(1):12, feb 2020.
- Liu, S., Xu, L., Zhao, Z., Erdélyi, R., Korsós, M. B., and Huang, X. Deep learning based solar flare forecasting model. II. Influence of image resolution. *The Astrophysical Journal*, 941(1):20, December 2022.
- Mason, J. P. and Hoeksema, J. T. Testing automated solar flare forecasting with 13 years of Michelson Doppler Imager magnetograms. *The Astrophysical Journal*, 723(1):634–640, November 2010.
- Moreland, K., Dayeh, M., Dayeh, M. A., Chatterjee, S., Bain, H., and Munoz-Jaramillo, A. A machine-learning oriented dataset for forecasting SEP occurrence and properties. In *44th COSPAR Scientific Assembly*, volume 44, pp. 1151, July 2022.
- Nishizuka, N., Sugiura, K., Kubo, Y., Den, M., and Ishii, M. Deep flare net (DeFN) model for solar flare prediction. *The Astrophysical Journal*, 858(2):113, May 2018.

- Panos, B. and Kleint, L. Real-time flare prediction based on distinctions between flaring and non-flaring active region spectra. *The Astrophysical Journal*, 891(1):17, feb 2020. URL <https://dx.doi.org/10.3847/1538-4357/ab700b>.
- Pesnell, W. D. Solar Dynamics Observatory (SDO). In *Handbook of Cosmic Hazards and Planetary Defense*, pp. 179–196. Springer, Cham, 2015.
- Qahwaji, R., Colak, T., Al-Omari, M., and Ipson, S. Automated prediction of CMEs using machine learning of CME - flare associations. *Solar Physics*, 248(2):471–483, April 2008.
- Sager, T. W. and Thisted, R. A. Maximum likelihood estimation of isotonic modal regression. *The Annals of Statistics*, 10(3):690 – 707, 1982. URL <https://doi.org/10.1214/aos/1176345865>.
- Sun, P., Dai, W., Ding, W., Feng, S., Cui, Y., Liang, B., Dong, Z., and Yang, Y. Solar flare forecast using 3D convolutional neural networks. *The Astrophysical Journal*, 941(1):1, December 2022.
- Sun, Z., Bobra, M. G., Wang, X., Wang, Y., Sun, H., Gombosi, T., Chen, Y., and Hero, A. Predicting solar flares using CNN and LSTM on two solar cycles of active region data. *The Astrophysical Journal*, 931(2):163, jun 2022. URL <https://dx.doi.org/10.3847/1538-4357/ac64a6>.
- Tang, R., Liao, W., Chen, Z., Zeng, X., song Wang, J., Luo, B., Chen, Y., Cui, Y., Zhou, M., Deng, X., Li, H., Yuan, K., Hong, S., and Wu, Z. Solar flare prediction based on the fusion of multiple deep-learning models. *The Astrophysical Journal Supplement Series*, 257(2):50, dec 2021. URL <https://dx.doi.org/10.3847/1538-4365/ac249e>.
- The SunPy Community, Mumford, S. J., Christe, S., Pérez-Suárez, D., Ireland, J., Shih, A. Y., Inglis, A. R., Liedtke, S., Hewett, R. J., Mayer, F., Hughitt, K., Freij, N., Meszaros, T., Bennett, S. M., Malocha, M., Evans, J., Agrawal, A., Leonard, A. J., Robitaille, T. P., Mampaey, B., Iván Campos-Rozo, J., and Kirk, M. S. SunPy—Python for solar physics. *Computational Science and Discovery*, 8(1):014009, January 2015.
- Tian, Y. Relationship between the evolution of solar magnetic field and flares in different active regions. *Journal of Physics: Conference Series*, 2282(1):012025, may 2022. URL <https://dx.doi.org/10.1088/1742-6596/2282/1/012025>.
- Um, T. T., Pfister, F. M. J., Pichler, D., Endo, S., Lang, M., Hirche, S., Fietzek, U., and Kulić, D. Data augmentation of wearable sensor data for Parkinson’s disease monitoring using convolutional neural networks. In *Proceedings of the 19th Association for Computing Machinery International Conference on Multimodal Interaction*, 2017.
- van Driel-Gesztelyi, L. and Green, L. M. Evolution of active regions. *Living Reviews in Solar Physics*, 12(1):1, Sep 2015. URL <https://link.springer.com/article/10.1007/lrsp-2015-1>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, 2017.
- Wang, X., Chen, Y., Toth, G., Manchester, W. B., Gombosi, T. I., Hero, A. O., Jiao, Z., Sun, H., Jin, M., and Liu, Y. Predicting solar flares with machine learning: Investigating solar cycle dependence. *The Astrophysical Journal*, 895(1):3, may 2020. URL <https://dx.doi.org/10.3847/1538-4357/ab89ac>.
- Wilks, D. S. Sampling distributions of the Brier score and Brier skill score under serial dependence. *Quarterly Journal of the Royal Meteorological Society*, 136(653):2109–2118, October 2010.
- Yu, D., Huang, X., Wang, H., and Cui, Y. Short-term solar flare prediction using a sequential supervised learning method. *Solar Physics*, 255(1):91–105, March 2009.
- Zerveas, G., Jayaraman, S., Patel, D., Bhamidipaty, A., and Eickhoff, C. A transformer-based framework for multivariate time series representation learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021.
- Zhang, H., Li, Q., Yang, Y., Jing, J., Wang, J. T. L., Wang, H., and Shang, Z. Solar flare index prediction using SDO/HMI vector magnetic data products with statistical and machine-learning methods. *The Astrophysical Journal Supplement Series*, 263(2):28, December 2022.