

Deep Learning for Active Region Classification: A Systematic Study from Convolutional Neural Networks to Vision Transformers

Edoardo Legnaro¹, Sabrina Guastavino^{1,2}, Michele Piana^{1,2}, and Anna Maria Massone¹

¹ MIDA, Dipartimento di Matematica, Università di Genova, via Dodecaneso 35 16146 Genova, Italy

² Istituto Nazionale di Astrofisica, Osservatorio Astrofisico di Torino, via Osservatorio 20 10025 Pino Torinese Italy

Abstract

A solar active region can significantly disrupt the Sun–Earth space environment, often leading to severe space weather events such as solar flares and coronal mass ejections. As a consequence, the automatic classification of active region groups is the crucial starting point for accurately and promptly predicting solar activity. This study presents our results concerned with the application of deep learning techniques to the classification of active region cutouts based on the Mount Wilson classification scheme. Specifically, we have explored the latest advancements in image classification architectures, from Convolutional Neural Networks to Vision Transformers, and reported on their performances for the active region classification task, showing that the crucial point for their effectiveness consists in a robust training process based on the latest advances in the field.

key words. Active region classification; deep learning; transformers

1 Introduction

The Sun is the source of a variety of dynamic phenomena, such as solar flares and coronal mass ejections (CMEs), which significantly impact space weather. These solar activities can cause disturbances that affect satellite operations, power grids, and communication systems, potentially leading to significant economic and technological disruptions (Baker 2002; Buzulukova & Tsurutani 2022). Accurate forecasting of space weather events is therefore crucial to mitigate their adverse effects on modern infrastructure (Guastavino et al. 2022, 2023; Georgoulis et al. 2024).

It is well-established that solar flares, i.e. intense bursts of radiation originating in the solar atmosphere, predominantly above sunspots (Hudson 2011; Benz 2017; Piana et al. 2022), are the main trigger of space weather. Therefore, the prediction of solar flares has become one of the most fundamental tasks of space weather understanding, and the development of data-driven models relying on artificial intelligence (AI) approaches is disruptively changing the scenario of this scientific and technological challenge (Guastavino et al. 2022, 2023). Recent studies have shown a strong correlation between the likelihood of solar flare occurrences and specific characteristics of active region groups, including their size, the number of sunspots, and their magnetic classification (McCloskey et al. 2018; Eren et al. 2016). It follows that active region classification represents the crucial initial step in predicting space weather.

Traditionally, expert observers manually classify active regions by examining images of the Sun, identifying sunspots, and categorizing them based on their size, shape and magnetic properties. This method has the disadvantages of being time-consuming, subjective, and not scalable with data volume. Nevertheless, the rapid accumulation of solar observations from an increasing number of space missions makes it possible to automating the classification of active region groups using machine learning techniques (Colak & Qahwaji 2008; Abd et al. 2010; Pesnell et al. 2012; Scherrer et al. 2012), and recent advancements in the field, particularly deep learning (LeCun et al. 2015; Goodfellow 2016), have shown promising results in this task (Fang et al. 2019; Tang et al. 2021). This study builds upon the research presented in (Fang et al. 2019) and (Tang et al. 2021) to explore and critically compare various state-of-the-art deep learning architectures for image classification, specifically focusing on convolutional neural networks (CNNs) and Vision Transformers (ViTs) when applied to solar active region cutouts. Specifically, we evaluated models trained on magnetograms, continuum images, or both, using 2D convolutions, and progressively trained models of increasing complexity to determine which architecture is most robust for this classification task.




Figure 1: Histogram displaying the distribution of aspect ratios (width/height) of the images in the SOLAR-STORM1 dataset.

This evaluation incorporates extensive on-the-fly data augmentation and presents results based on five-fold cross-validation.

The plan of the paper is as follows. Section 2 describes the properties of the data set used for the experiments, and the corresponding pre-processing step. Section 3 provides details on the design and properties of the deep learning networks employed for the analysis. Section 4 describes the results of the study. Our conclusions are offered in Section 5.

2 Data

2.1 The Dataset

This study considered the SOLAR-STORM1 dataset provided by the Space Environment Warning and AI Technology Interdisciplinary Innovation Working Group (Fang et al. 2019)¹. This dataset contains continuous and magnetogram images provided by the Spaceweather HMI Active Region Patch (SHARP) from the Helioseismic and Magnetic Imager (HMI) on-board the Solar Dynamics Observatory (Pesnell et al. 2012; Scherrer et al. 2012). The images in this dataset are sunspots cutouts of different aspect ratios, as shown in Figure 1.

The temporal range of the dataset covers the interval from 2010 May 1 to 2017 December 12 (see Figure 2, top panel) and the classes distribution is shown in Table 1, second column. Specifically, the Mount Wilson sunspot magnetic types are divided into the following 8 classes.

- Alpha: Unipolar sunspot group.
- Beta: Bipolar sunspot group with distinct positive and negative magnetic polarities.
- Gamma: Complex region with irregular positive and negative polarities, preventing bipolar classification.
- Delta: Umbrae with opposite polarities within one penumbra, separated by less than 2° .
- Beta-Gamma: Bipolar group with complex polarity distribution, lacking a clear dividing line.
- Beta-Delta: Beta group containing one or more Delta spots.
- Beta-Gamma-Delta: Beta-Gamma group containing one or more Delta spots.
- Gamma-Delta: Gamma group containing one or more Delta spots.

Due to the limited number of examples in some groups, the Delta and Gamma-Delta magnetic types were excluded, and the Beta-Gamma, Beta-Delta, and Beta-Gamma-Delta classes were merged into a single class, referred to as Beta-X. As a result, the problem was approached as a multi-class classification task with three categories: Alpha, Beta, and Beta-X. In this dataset, images were taken every 96 minutes,

¹The dataset is available at <https://tianchi.aliyun.com/dataset/dataDetail?dataId=74779>

Class	Full Dataset Counts (%)	Test Dataset Counts (%)
Alpha	5276 (33.73%)	567 (48.38%)
Beta	7849 (50.18%)	496 (42.32%)
Beta-X	2516 (16.09%)	109 (9.30%)
Total	15641	1172

Table 1: Distribution of class cardinality and rate in the whole SOLAR-STORM1 dataset and in the test set.

	Train				Val			
	Alpha (%)	Beta (%)	Beta-X (%)	Total	Alpha (%)	Beta (%)	Beta-X (%)	Total
Fold 1	3821 (33.01%)	5691 (49.17%)	2063 (17.82%)	11575	888 (30.68%)	1662 (57.43%)	344 (11.89%)	2894
Fold 2	3845 (33.22%)	5547 (47.92%)	2183 (18.86%)	11575	864 (29.85%)	1806 (62.40%)	224 (7.74%)	2894
Fold 3	3992 (34.49%)	5849 (50.53%)	1734 (14.98%)	11575	717 (24.78%)	1504 (51.97%)	673 (23.26%)	2894
Fold 4	3516 (30.38%)	6067 (52.41%)	1992 (17.21%)	11575	1193 (41.22%)	1286 (44.44%)	415 (14.34%)	2894
Fold 5	3662 (31.63%)	6258 (54.06%)	1656 (14.31%)	11576	1047 (36.19%)	1095 (37.85%)	751 (25.96%)	2893

Table 2: Distribution of class cardinality and rate in the training and validation sets across folds.

but a sunspot group does not change consistently within 24 hours (see McCloskey et al. (2016)). Therefore, this dataset presents a high degree of redundancy, and a random split of the training, validation, and test sets (as done in Fang et al. (2019)) would result in inflated test scores. As the dataset lacks active region identifiers, such as NOAA numbers, the only viable option was to apply a temporal split to ensure that highly similar images are not distributed across different sets. This approach mirrors the splitting strategy used in Tang et al. (2021), where all data after 2016 January 1 were designated as the test set, and the remaining data were used for fivefold cross-validation. A visualization of this splitting method is presented at the bottom of Figure 2. The number of classes along with their respective percentages for the training and validation sets across the five folds are detailed in Table 2, while the distribution for the test set is shown in Table 1, third column.

2.2 Data pre-processing

Data augmentation is the process of artificially generating new data from existing one, primarily to train deep learning models, which require a large amount of training data samples in order to avoid overfitting. In this work the images are first normalized and cropped following the approach described in Tang et al. (2021). In particular, magnetograms were normalized so that each pixel m_{out} of the normalized image follows from

$$m_{out} = \text{HardTanh}\left(\frac{m_{in}}{d}\right), \quad (1)$$

where m_{in} represents the original pixel value and $d \geq 0$. The HardTanh function is defined as:

$$\text{HardTanh}(x) = \begin{cases} 1 & x \geq 1 \\ x & -1 < x < 1 \\ -1 & x \leq -1 \end{cases}, \quad (2)$$

which ensures that the final value lies within the interval $[-1, 1]$. In (1) we set $d = 800$, as done in Fang et al. (2019) and Tang et al. (2021), since this threshold value appears to give slightly higher performance metrics compared to higher constant values (see Table 3). For continuum images, normalization is performed such that each pixel c_{out} is determined by

$$c_{out} = 1 - \frac{c_{in} - \min}{\max - \min}, \quad (3)$$

where min and max respectively denote the minimum and maximum values of the input image, which has pixels c_{in} . The final normalized values lie in the interval $[0, 1]$, with background closer to 0 and the sunspots closer to 1. Additionally, we included cropped images of the active regions in the dataset, using a strategy similar to that of Tang et al. (2021). Specifically, this cropping was performed by comparing the absolute differences between the maximum values in the rows and columns of the normalized continuum images.

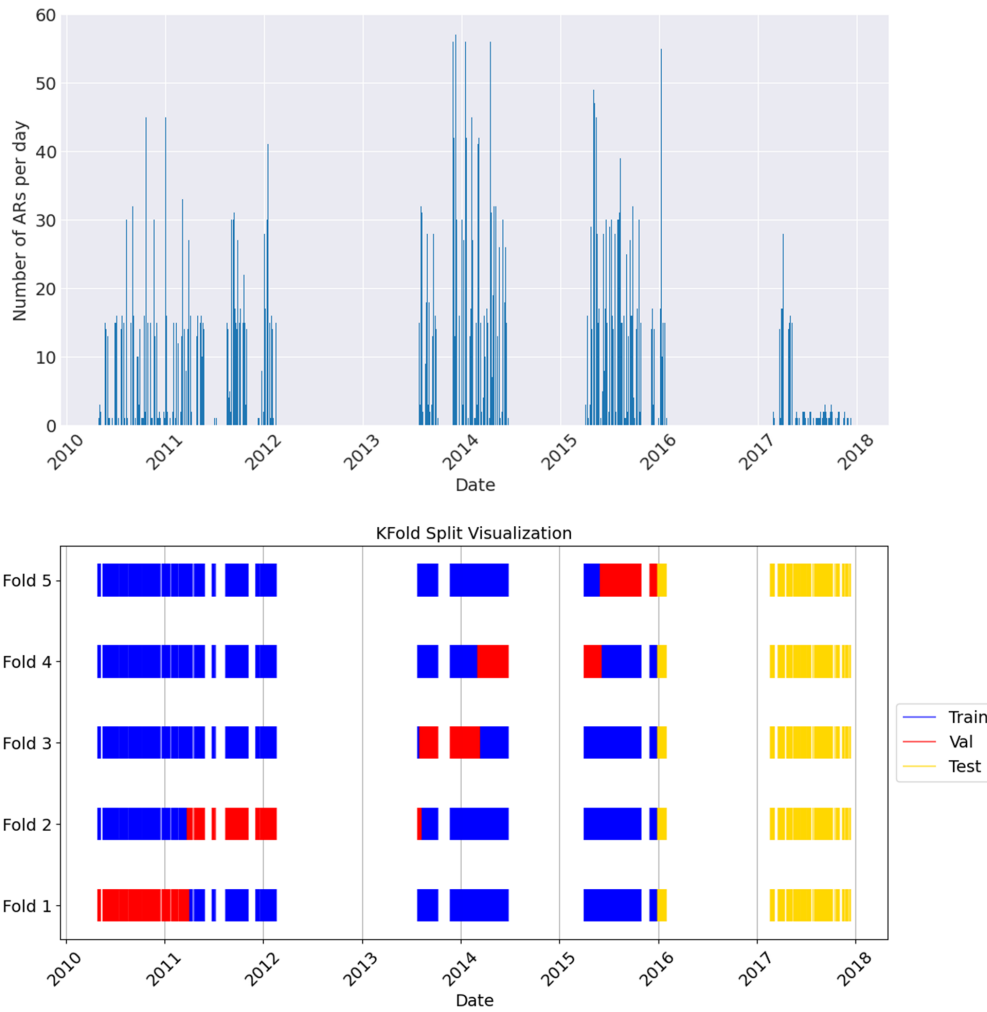


Figure 2: Temporal distribution of Active Regions (ARs) and the K-Fold split visualization for the dataset. The upper panel shows the number of ARs per day present in the dataset, while the lower one illustrates how the data is divided into training (blue), validation (red), and test (yellow) sets across the five different folds.




Figure 3: Effects of the normalization process on the images. The AR considered as an example is a Beta-X one at time 2014-03-17 16:00. The first row shows the original continuum and magnetogram images, while the second one shows, on the left, the normalized continuum image with Min-Max normalization (3) and, on the right, the magnetogram image normalized using the HardTanh function (1) with constant $d = 800$. The last line shows the cropping of the AR done to filter out noise and allow the model to concentrate on the significant features relevant to the classification task.

Peaks in these differences were then identified, and a bounding box was selected by adding a margin around these peaks. This is done in order to filter out noise, allowing the model to concentrate on the significant features relevant to the classification task. An example of the normalization and cropping is shown in Figure 3. Next, we apply a random horizontal and/or vertical flip of the images, a perspective change, and an affine transformation including a minor translation and a rotation of a random angle up to 30° . Usually, data augmentation techniques are applied before fitting a deep learning model on the whole dataset. However, in this way a single augmented dataset is created. Performing data augmentation on-the-fly as in Cerqueira et al. (2024), where a new set of random transformations is applied to the dataset at each epoch, creates slight variations with every iteration. Therefore, this approach enables better exploration of the data space and can help reduce the risk of overfitting.

3 Deep neural networks

In this work, we explored different neural network architectures for classifying sunspots by their magnetic classes. To achieve this, we utilized the models provided in Wightman (2019). The neural network architectures selected for this comparative study encompass both Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs), since traditionally CNNs are better for capturing local features, while transformers focus more on global features. By comparing these two distinct architectures, this study aimed to

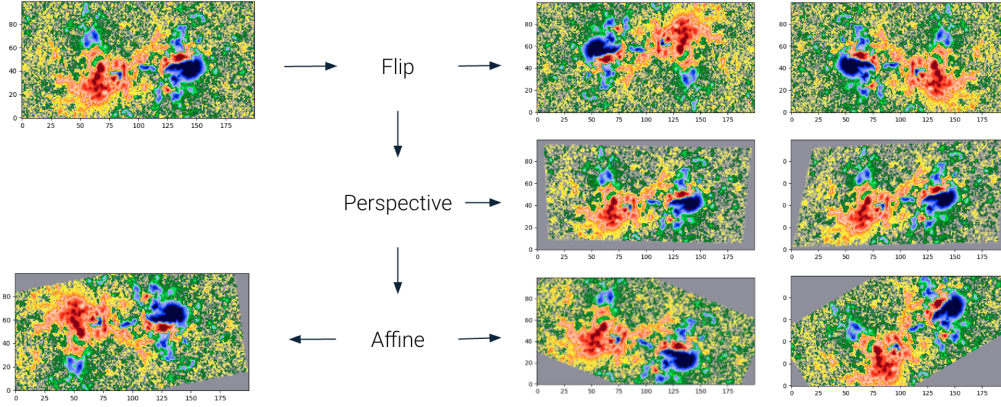


Figure 4: Illustration of data augmentation techniques. The original image (top left) undergoes a series of transformations to enhance the model’s ability to generalize. The first row shows the effect of horizontal and vertical flips. The second row demonstrates perspective transformations, altering the viewing angle. The third row applies affine transformations, including scaling, rotation, and shear, to further diversify the dataset.

d	Acc. (min - max)	F1 Score (min - max)	F1 Alpha (min - max)	F1 Beta (min - max)	F1 Beta-X (min - max)
800	0.8563 (0.8396 - 0.8805)	0.8223 (0.7972 - 0.8472)	0.9140 (0.8994 - 0.9295)	0.8215 (0.7956 - 0.8551)	0.7315 (0.6829 - 0.7570)
1200	0.8466 (0.8191 - 0.8814)	0.8048 (0.7659 - 0.8441)	0.9159 (0.9103 - 0.9335)	0.8089 (0.7750 - 0.8595)	0.6895 (0.6115 - 0.7451)
2000	0.8398 (0.8276 - 0.8532)	0.7961 (0.7800 - 0.8079)	0.9117 (0.8868 - 0.9241)	0.7963 (0.7712 - 0.8161)	0.6804 (0.6241 - 0.7424)
4000	0.8493 (0.8106 - 0.8763)	0.8104 (0.7657 - 0.8477)	0.9135 (0.8881 - 0.9275)	0.8111 (0.7641 - 0.8496)	0.7066 (0.6449 - 0.7717)

Table 3: Performance metrics (Accuracy, Macro F1 Score (macro-averaged) and F1 Scores for the singles classe Alpha, Beta and Beta-X) for different divisor values (d). The values are presented as the average, minimum and maximum value across the five folds for each metric obtained on the test set. The model trained here is resnet18 on magnetograms, using the Adam optimizer with learning rate 10^{-5} and images resized to 224x224 squares.

evaluate their respective strengths and limitations in classifying solar active region cutouts, with the goal of determining the most robust approach for this classification task.

3.1 Convolutional Neural Networks

A Convolutional Neural Network (CNN) is a specific type of deep learning architecture designed for efficiently processing grid-like data, such as images. CNNs use convolutional layers to automatically extract hierarchical features by applying learned filters across the input data. Pooling layers further reduce spatial dimensions, helping to downsample feature maps and minimize computational complexity. Non-linear activation functions then introduce complexity into the model, allowing it to capture intricate data patterns. Finally, fully connected layers transform the learned features into a final prediction. CNNs excel at visual tasks due to their ability to learn spatial relationships and patterns directly from data (LeCun et al. 1998; Krizhevsky et al. 2012). We now briefly introduce the CNNs whose performances have been explored in this paper.

3.1.1 VGG

The VGG (Visual Geometry Group) architecture, introduced in Simonyan & Zisserman (2014), consist in of sequential layers of small 3x3 convolutions stacked together, followed by max-pooling layers. This design allows the network to capture increasingly complex features while maintaining manageable computational complexity.

3.1.2 Inception

The Inception network, introduced by Szegedy et al. (2015) and also known as GoogLeNet, is a deep convolutional neural network involving the “Inception module,” which applies multiple convolutional filters of varying sizes (1x1, 3x3, and 5x5), and a pooling layer in parallel. This allows the network to capture features at different scales without significantly increasing the computational cost. Additionally, the architecture incorporates dimensionality reduction through 1x1 convolutions, which helps controlling the number of parameters and memory usage.

3.1.3 ResNet

The ResNet (Residual Network) architecture, introduced in He et al. (2016), is a type of deep neural network that introduces connections or residual connections to address the vanishing gradient problem occurring in very deep neural networks. These connections bypass one or more layers by allowing the input of a layer to be added directly to its output, in such a way that the network passes information directly to deeper layers. This technique enables the model to learn identity mappings, making it easier to train deeper networks without degradation in performance.

3.2 Vision Transformers

Visual Transformers (ViTs) are a class of deep learning architectures designed to handle visual data using the transformer model, which was originally developed for natural language processing tasks (Dosovitskiy 2020). Unlike CNNs, ViTs operate on image patches that are treated as a sequence of tokens, similar to words in a sentence, and are processed using self-attention mechanisms (see Vaswani (2017)). This mechanism is used to determine the importance of each element in a sequence relative to others. It allows the model to focus on relevant parts of the input data when generating a representation for each element. In self-attention, for each input token (or element), the model computes three vectors: query, key, and value. The query from one token is compared with the keys of all other tokens to compute attention scores, which indicate how much focus the model should place on each token. These scores are then used to compute a weighted sum of the value vectors, creating a new representation for each token based on its relationships to others. In this way, the model can capture long-range dependencies within the data, as each token can attend to all others, regardless of their distance in the input sequence. This approach allows ViTs to capture global dependencies between patches early in the network, making them effective for tasks requiring global context.

Besides ViTs architectures based on Dosovitskiy (2020), in this paper we also explore the following variations.

3.2.1 DeiT

The Data-efficient Image Transformer (DeiT), introduced by (Touvron et al. 2021), is an improvement on ViT that focuses on reducing the data requirements for training Vision Transformers. While ViT performs well with large datasets, DeITs introduce techniques such as knowledge distillation to train the model more efficiently on smaller datasets without sacrificing performances. DeITs retain the transformer-based architecture of ViTs but optimize their training process, making it more accessible for practical applications where large amounts of labeled data may not be available.

3.2.2 BEiT

The Bidirectional Encoder representation from Image Transformers (BEiT), introduced by (Bao et al. 2021), is a self-supervised vision transformer model that extends the principles of BERT-style pre-training from natural language processing to vision tasks. BEiT leverages masked image modeling, where portions of the input image patches are masked and the model is trained to predict the missing content based on the surrounding context. This pre-training technique allows BEiT to learn rich image representations without the need for extensive labeled data, making it highly effective in scenarios with limited supervision. After pre-training, BEiT can be fine-tuned for various downstream tasks, such as image classification, object detection, and segmentation. BEiT’s success lies in its ability to transfer knowledge learned during self-supervised pre-training to a wide range of vision applications, achieving competitive results while reducing the reliance on large labeled datasets.

4 Results

The results presented in this section have been obtained by applying the pipeline summarized in Figure 5 to all the networks previously illustrated. To train these networks, we used the Adam optimizer (Kingma & Ba 2015) with a learning rate of 10^{-5} , employing a weighted categorical cross-entropy loss function to address class imbalance in the different folds. Once normalized, the images have been resized to a 224×224 square, with zero padding added to maintain the original aspect ratio. This was done so that the weights of the neural networks are initialized using transfer learning from pre-trained models on the ImageNet dataset (Deng et al. 2009). This approach leverages the feature extraction capabilities learned from a large and diverse set of images, allowing the models to start with a strong foundation rather than from scratch. Even though transfer learning performs better when the source and target domains are closely related, it still offers significant advantages by improving convergence and potential accuracy. Additionally, early stopping was implemented to halt training if the validation loss did not improve. The model weights were then selected based on the epoch that achieved the best validation performance. To enhance the statistical reliability, the results are presented as the average test scores across all folds, along with their corresponding minimum and maximum values. Models have been trained on magnetogram data, continuum data and both combined in a 2 channel image using 2D convolutions. The number of trainable parameters for each model is reported in Table 4. The results of the analysis are summarized in Table 5.

Among all the models, ResNet18 trained on both magnetogram and continuum images achieved the highest average accuracy (0.903) and highest average F1 score (0.864), demonstrating the efficacy of combining both image types in providing richer information for model training. This suggests that the ResNet architecture can effectively leverage diverse inputs to outperform other architectures in terms of overall classification accuracy.

When analyzing models trained on only magnetograms, the DEiT Base model stands out, achieving the highest average accuracy (0.901) and highest F1 score of 0.861. This indicates that transformer-based architectures can capture relevant spatial features present in magnetograms (even if ResNet-based architectures remain competitive also in this case). For models trained exclusively on continuum images, ResNet10t achieved the highest average accuracy (0.877) and ResNet18 the highest average F1 score (0.825). This suggests that continuum images may contain less spatial variability than magnetograms, and a simple ResNet-based architecture with relatively few parameters is sufficient to extract meaningful patterns from the data.

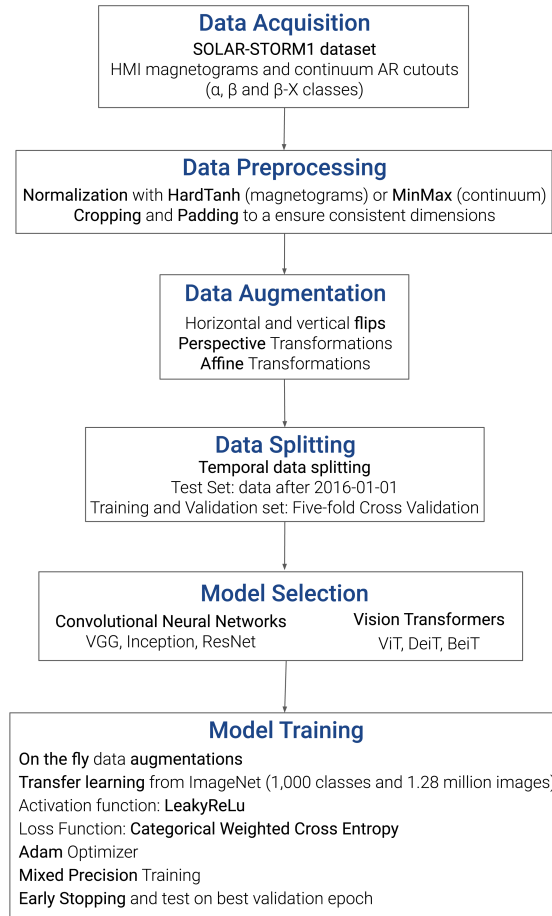


Figure 5: A summary of the workflow followed in this study.

Model	Trainable Parameters (Millions)
ResNet10t	4.92
ResNet18	11.2
ResNet26	14.9
Inception V4	41.1
VGG11	128.8
ViT Base	85.7
DEiT Base	85.7
BEiT Base	85.7

Table 4: Approximate value of trainable parameters for the models considered in this study. The number of trainable parameters is computed by iterating over all parameters in each model and summing the number of elements involved in the computation of gradients during backpropagation. The models trained from the timm library (Wightman 2019) are: resnet10t, resnet18, resnet26, inception_v4, vgg11, vit_base_patch16_224, deit_base_patch16_224, beit_base_patch16_224.

Model	image type	Acc. (min-max- Δ)	F1 Score (min-max- Δ)	F1 α (min-max- Δ)	F1 β (min-max- Δ)	F1 β -X (min-max- Δ)
ResNet10t	magnetogram	0.893 (0.871-0.904-0.033)	0.849 (0.837-0.863-0.026)	0.949 (0.918-0.963-0.045)	0.872 (0.852-0.885-0.033)	0.726 (0.697-0.748-0.051)
	continuum	0.877 (0.845-0.893-0.049)	0.821 (0.779-0.838-0.059)	0.956 (0.946-0.961-0.015)	0.848 (0.797-0.872-0.074)	0.659 (0.586-0.686-0.100)
ResNet18	both	0.894 (0.870-0.916-0.046)	0.847 (0.816-0.877-0.061)	0.953 (0.940-0.963-0.023)	0.874 (0.849-0.902-0.054)	0.715 (0.653-0.769-0.116)
	magnetogram	0.887 (0.876-0.897-0.020)	0.843 (0.825-0.858-0.033)	0.949 (0.941-0.957-0.016)	0.863 (0.845-0.876-0.031)	0.718 (0.676-0.759-0.083)
ResNet26	continuum	0.872 (0.853-0.887-0.033)	0.825 (0.795-0.840-0.045)	0.943 (0.925-0.949-0.025)	0.842 (0.808-0.863-0.055)	0.689 (0.628-0.709-0.081)
	both	0.903 (0.887-0.921-0.033)	0.864 (0.845-0.891-0.046)	0.953 (0.941-0.960-0.019)	0.885 (0.868-0.906-0.038)	0.756 (0.722-0.812-0.090)
Inception V4	magnetogram	0.882 (0.844-0.909-0.065)	0.841 (0.800-0.883-0.083)	0.940 (0.915-0.956-0.041)	0.859 (0.817-0.895-0.078)	0.724 (0.667-0.812-0.145)
	continuum	0.865 (0.817-0.904-0.087)	0.823 (0.777-0.869-0.092)	0.925 (0.888-0.949-0.061)	0.840 (0.782-0.886-0.103)	0.705 (0.662-0.774-0.112)
VGG11	both	0.883 (0.852-0.916-0.064)	0.837 (0.788-0.881-0.092)	0.945 (0.934-0.959-0.025)	0.861 (0.823-0.900-0.077)	0.704 (0.606-0.782-0.176)
	magnetogram	0.897 (0.883-0.903-0.020)	0.856 (0.843-0.864-0.021)	0.948 (0.935-0.961-0.027)	0.880 (0.866-0.887-0.021)	0.739 (0.722-0.758-0.036)
ViT Base	continuum	0.866 (0.838-0.892-0.054)	0.817 (0.780-0.851-0.070)	0.937 (0.924-0.948-0.023)	0.838 (0.799-0.872-0.073)	0.677 (0.608-0.735-0.127)
	both	0.854 (0.839-0.863-0.025)	0.800 (0.773-0.819-0.046)	0.919 (0.904-0.947-0.044)	0.832 (0.812-0.849-0.036)	0.649 (0.584-0.702-0.118)
DEiT Base	magnetogram	0.894 (0.868-0.925-0.057)	0.843 (0.807-0.882-0.076)	0.950 (0.922-0.966-0.044)	0.875 (0.853-0.912-0.059)	0.706 (0.644-0.769-0.126)
	continuum	0.875 (0.840-0.896-0.055)	0.829 (0.798-0.850-0.053)	0.938 (0.907-0.958-0.052)	0.851 (0.816-0.874-0.059)	0.696 (0.672-0.724-0.052)
BEiT Base	both	0.890 (0.858-0.910-0.051)	0.832 (0.794-0.871-0.077)	0.947 (0.933-0.959-0.026)	0.873 (0.833-0.896-0.062)	0.677 (0.583-0.750-0.099)
	magnetogram	0.882 (0.871-0.896-0.025)	0.835 (0.814-0.855-0.042)	0.942 (0.916-0.956-0.040)	0.859 (0.834-0.875-0.042)	0.704 (0.651-0.750-0.099)
DEiT Base	continuum	0.848 (0.794-0.905-0.112)	0.803 (0.740-0.866-0.126)	0.919 (0.875-0.953-0.078)	0.817 (0.744-0.885-0.141)	0.673 (0.571-0.761-0.190)
	both	0.889 (0.863-0.915-0.052)	0.845 (0.803-0.869-0.066)	0.945 (0.918-0.959-0.040)	0.867 (0.817-0.903-0.086)	0.723 (0.633-0.765-0.133)
BEiT Base	magnetogram	0.901 (0.885-0.916-0.031)	0.861 (0.840-0.885-0.045)	0.952 (0.933-0.968-0.035)	0.881 (0.856-0.901-0.045)	0.749 (0.700-0.802-0.102)
	continuum	0.864 (0.834-0.892-0.059)	0.817 (0.778-0.848-0.070)	0.932 (0.911-0.951-0.041)	0.838 (0.789-0.872-0.083)	0.680 (0.611-0.723-0.112)
BEiT Base	both	0.902 (0.876-0.926-0.049)	0.853 (0.812-0.888-0.076)	0.957 (0.925-0.968-0.043)	0.884 (0.865-0.912-0.048)	0.718 (0.645-0.786-0.142)
	magnetogram	0.873 (0.854-0.895-0.041)	0.828 (0.801-0.856-0.055)	0.937 (0.917-0.950-0.033)	0.846 (0.819-0.874-0.055)	0.702 (0.646-0.748-0.102)
BEiT Base	continuum	0.866 (0.842-0.902-0.060)	0.817 (0.786-0.860-0.074)	0.938 (0.918-0.952-0.034)	0.838 (0.801-0.883-0.081)	0.676 (0.616-0.745-0.129)
	both	0.891 (0.844-0.920-0.076)	0.845 (0.772-0.886-0.114)	0.942 (0.907-0.961-0.054)	0.873 (0.831-0.904-0.073)	0.720 (0.577-0.793-0.216)

Table 5: Performance results for different deep learning models. The table reports the mean values of the metrics across the five folds, along with minimum and maximum values and the difference Δ between maximum and minimum values. Model names are the same as in Table 4.

In general, accuracy on continuum images tends to be slightly lower across all models compared to magnetograms or combined image types, likely due to the more homogeneous nature of this data type.

When analyzing the F1 score for the α class, ResNet10t and ResNet18 achieve the highest average scores (0.953) when trained on both image types, while for the β class the highest F1 score on average is obtained with the ResNet18 network. Models trained only on magnetograms, such as ResNet18 and DEiT Base, also performed well, achieving F1 β scores of 0.876 and 0.881, respectively. However, models trained solely on continuum data, like ViT Base (F1 β = 0.817), showed a noticeable drop in performance, possibly due to the limited variability in continuum images. In terms of the F1 β -X score, which evaluates performance in the most challenging class, as expected the DEiT architecture achieves the best average score (0.749), suggesting that transformer-based models are better suited for capturing the intricate spatial relationships present in complex magnetic structures. However, performance declines notably when models are trained exclusively on continuum images.

The observed variability across folds, indicated by Δ values, generally ranges from 2% to 10% for accuracy and overall F1 scores. For the more complex β -X class, the variation can be even higher, with differences of up to 20%. Anyway, in general, the average performances do not change drastically from one model architecture to another. This could be attributed to several factors. First, complex models may not fully utilize their capacity if the dataset does not present sufficient complexity or variability, leading to comparable performances with simpler architectures that can efficiently capture the essential patterns. Additionally, simpler models benefit from reduced risk of overfitting due to their smaller number of parameters, particularly when trained on datasets with moderate complexity. This is further enhanced by the use of modern training techniques such as batch normalization, dropout, and advanced optimizers, which improve generalization across models. These factors collectively explain why more complex architectures may not consistently outperform simpler ones, particularly if a specific dataset does not demand the added complexity. This closeness can also in part be attributed to the extensive data augmentation procedure employed during the training phase. In particular, ViTs tend to be less sensitive to data augmentation compared to traditional CNNs, mainly due to their global attention mechanism and ability to model relationships across the entire image. Transformers inherently capture more comprehensive image representations, making them more robust to small variations that data augmentations typically address in CNNs. As a result, extensive data augmentations help to narrow the performance gap between CNNs and more complex transformer models.

Compared to similar works, performance metrics reported by the study of Fang et al. (2019) are not directly comparable due to the random splitting of a highly redundant dataset. Comparison with the study of Tang et al. (2021), where 3D convolutions are used to classify an image consisting of both magnetogram and continuum data, shows that the maximum scores obtained with 2D convolutions are comparable to a more involved ensemble model. However, our goal was not to focus on the individual absolute values of the scores in comparison to other works, as it is possible to continue training with different random seeds in search of a favourable outlier that yields higher scores (Picard 2021). Instead, we aimed to provide a meaningful comparison of the performance of different model architectures under similar conditions,

highlighting their relative strengths and weaknesses from a global perspective.

5 Conclusions

In this study, we analyzed different deep learning architectures for the task of classifying the magnetic class of solar cutouts. We observed that combining magnetogram and continuum image types enhances model performance by leveraging complementary features from diverse inputs. When considering only magnetograms, data efficient transformer models achieve the best performance, demonstrating the effectiveness of transformer-based architectures in capturing the spatial complexity of magnetograms. Models trained exclusively on continuum images exhibit overall lower performance, suggesting that continuum images, due to their more homogeneous nature, offer less spatial variability.

Overall, performance metrics do not change drastically between model architectures, a fact that can be attributed to several factors: complex models may not fully utilize their capacity if the dataset lacks sufficient complexity or variability, resulting in performance comparable to simpler architectures that can effectively capture the essential patterns; simpler models have a reduced risk of overfitting, particularly when trained on moderately complex datasets; modern training techniques such as batch normalization, dropout, and advanced optimizers further improve generalization, helping to reduce the performance gap between simpler and more complex models. Finally, the extensive data augmentation procedures employed during the training phase also contribute to closing the gap between simpler CNNs and more complex ViTs.

Acknowledgement

EL and AMM acknowledge the HORIZON Europe ARCAFF Project, Grant No. 101082164. SG, MP and AMM acknowledge INdAM-GNCS; SG was supported by the Programma Operativo Nazionale (PON) “Ricerca e Innovazione” 2014–2020.

References

- Abd, M. A., Majed, S. F., & Zharkova, V. 2010, in *Technological developments in networking, education and automation*, Springer, 321–325
- Baker, D. N. 2002, *Science*, 297, 1486
- Bao, H., Dong, L., Piao, S., & Wei, F. 2021, arXiv preprint arXiv:2106.08254
- Benz, A. O. 2017, *Living reviews in solar physics*, 14, 1
- Buzulukova, N., & Tsurutani, B. 2022, *Frontiers in Astronomy and Space Sciences*, 9, 1017103
- Cerqueira, V., Santos, M., Baghoussi, Y., & Soares, C. 2024, arXiv preprint arXiv:2404.16918
- Colak, T., & Qahwaji, R. 2008, *Solar Physics*, 248, 277
- Deng, J., Dong, W., Socher, R., et al. 2009, in *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 248–255
- Dosovitskiy, A. 2020, arXiv preprint arXiv:2010.11929
- Eren, S., Kilcik, A., Atay, T., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, stw2742
- Fang, Y., Cui, Y., Ao, X., et al. 2019, *Advances in Astronomy*
- Georgoulis, M. K., Yardley, S. L., Guerra, J. A., et al. 2024, *Advances in Space Research*
- Goodfellow, I. 2016, *Deep learning*, Vol. 196 (MIT press)
- Guastavino, S., Marchetti, F., Benvenuto, F., Campi, C., & Piana, M. 2022, *Astronomy & Astrophysics*, 662, A105
- . 2023, *Frontiers in Astronomy and Space Sciences*, 9,

- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in Proceedings of the IEEE conference on computer vision and pattern recognition, 770–778
- Hudson, H. S. 2011, *Space Science Reviews*, 158, 5
- Kingma, D. P., & Ba, J. 2015, in 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, ed. Y. Bengio & Y. LeCun.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, *Advances in neural information processing systems*, 25
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *nature*, 521, 436
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proceedings of the IEEE*, 86, 2278
- McCloskey, A. E., Gallagher, P. T., & Bloomfield, D. S. 2016, *Solar Physics*, 291, 1711
- . 2018, *Journal of Space Weather and Space Climate*, 8, A34
- Pesnell, W., Thompson, B., & Chamberlin, P. 2012, *The solar dynamics observatory (SDO)* (Springer)
- Piana, M., Emslie, A. G., Massone, A. M., & Dennis, B. R. 2022, *Hard X-Ray Imaging of Solar Flares*, Vol. 164 (Springer)
- Picard, D. 2021, arXiv preprint arXiv:2109.08203
- Scherrer, P. H., Schou, J., Bush, R., et al. 2012, *Solar Physics*, 275, 207
- Simonyan, K., & Zisserman, A. 2014, arXiv preprint arXiv:1409.1556
- Szegedy, C., Liu, W., Jia, Y., et al. 2015, in Proceedings of the IEEE conference on computer vision and pattern recognition, 1–9
- Tang, R., Zeng, X., Chen, Z., et al. 2021, *The Astrophysical Journal Supplement Series*, 257, 38
- Touvron, H., Cord, M., Douze, M., et al. 2021, in International conference on machine learning, PMLR, 10347–10357
- Vaswani, A. 2017, *Advances in Neural Information Processing Systems*
- Wightman, R. 2019, *PyTorch Image Models*,