

AI Foundation Model for Heliophysics: Applications, Design, and Implementation

Sujit Roy^{1,2*†}, Talwinder Singh^{5*†}, Marcus Freitag^{3†},
Johannes Schmude^{3†}, Rohit Lal^{1†}, Dinesha Hegde^{6†},
Soumya Ranjan⁴, Amy Lin¹, Vishal Gaur¹, Etienne Eben Vos³,
Rinki Ghosal¹, Badri Narayana Patro⁸, Berkay Aydin⁷,
Nikolai Pogorelov⁶, Juan Bernabe Moreno³, Manil Maskey²,
Rahul Ramachandran²

¹Earth System Science Center, The University of Alabama in
Huntsville, Huntsville, Alabama, USA.

²NASA Marshall Space Flight Center, Huntsville, Alabama, USA.

³ IBM Research.

⁴ Development Seed, Washington, DC, USA.

⁵ Department of Physics & Astronomy, Georgia State University.

⁶Center for Space Plasma and Aeronomic Research (CSPAR), The
University of Alabama in Huntsville, Huntsville, AL, USA.

⁷Department of Computer Science, Georgia State University, Atlanta,
United States.

⁸Microsoft Research, India.

*Corresponding author(s). E-mail(s): sujit.roy@nasa.gov,
tsingh14@gsu.edu;

†These authors contributed equally to this work.

Abstract

Deep learning-based methods have been widely researched in the areas of language and vision, demonstrating their capacity to understand long sequences of data and their usefulness in numerous helio-physics applications. Foundation models (FMs), which are pre-trained on a large-scale datasets, form the basis for a variety of downstream tasks. These models, especially those based on transformers in vision and language, show exceptional potential for adapting to a wide range of downstream applications. In this paper, we provide our perspective on

the criteria for designing a FM for heliophysics and associated challenges and applications using the Solar Dynamics Observatory (SDO) dataset. We believe that this is the first study to design a foundation model in the domain of heliophysics.

Keywords: Heliophysics, AI, Foundation Model, Downstream Applications

1 Introduction

Since the launch of the Solar Dynamics Observatory (SDO), it has been providing a continuous high-resolution observation of the Sun [1]. SDO has been a very fruitful mission with ≈ 6000 peer-reviewed papers published using its data [2]. SDO’s Helioseismic and Magnetic Imager [HMI; 3] provides high resolution (0.5 arcsec pixels) vector magnetograms and Doppler images, which are used to study the evolution of the magnetic field and plasma velocity on the solar surface respectively. The Atmospheric Imaging Assembly [AIA; 4] onboard SDO provides full-disk high-resolution (0.6 arcsec pixels) images of the Sun in seven extreme ultraviolet (EUV) wavelengths that can capture plasma in the solar atmosphere starting from the chromosphere to the upper corona with temperatures from $\approx 20,000$ K to more than 20 million K. AIA also images the Sun in two ultraviolet (UV) and one visible wavelength bands. AIA data have been instrumental in understanding the evolution of various coronal structures, such as coronal holes (CHs) [5, and references therein], active regions (ARs) [6, and references therein], coronal loops [7, and references therein], etc. This data has also led to better understanding of coronal processes such as EUV wave propagation [8, and references therein], solar flare eruptions [9, and references therein], filament eruptions [10, and references therein], etc. Figure 1 shows the line-of-sight (LOS) magnetogram captured by HMI in the top-left panel. Other panels show the Sun observed in seven EUV wavelengths by AIA. Wavelengths in \AA are denoted in the bottom right corners of all panels.

SDO has produced (and continues to produce) an immense amount of data during its lifetime, sending down nearly 1.5 terabytes of data daily. SDO’s large database provides a perfect opportunity to use machine learning (ML) based methods to solve many of the outstanding problems in solar and heliospheric physics. In this paper, we discuss how a class of ML models called Foundation Models (FMs) can be trained using the SDO data, detailing the implementation, planned downstream tasks, and preliminary results.

FMs generally consist of an encoder f_θ and a decoder g_ϕ , where θ and ϕ represent the model parameters. Pre-training the FM involves optimizing

$$g_\phi \circ f_\theta$$

using a self-supervised learning approach [11]. When fine-tuning for specific downstream tasks, the encoder f_θ is pretained, but the decoder g_ϕ is replaced by a




Fig. 1: Data captured by HMI and AIA instruments aboard SDO on 2012-01-30 at 22:12 UT. The top left panel shows the line-of-sight magnetogram from HMI and the rest of the panels show Sun in the seven EUV wavelengths from AIA. The wavelengths in Å are mentioned in the bottom right corner of each panel.

task-specific decoder h_ψ . However, if the fine tuning task involves very similar outcome like pretraining objective, we can keep the decoder and just change its head. In fine-tuning, the model then optimizes $h_\psi \circ f_\theta$. Typically, decoders are shallow, whereas encoders can have millions or even billions of parameters, as they are designed to understand and represent the input data’s complexity. This approach reduces the training burden by focusing on task-specific decoders. Furthermore, the pre-training data must be comprehensive enough that the learned distribution captures the variables and underlying principles necessary to generalize to various downstream tasks.

Recent studies [12–15] have focused on developing a versatile AI system capable of handling various tasks by building FMs for earth sciences. Initially introduced in the domains of natural language processing (NLP) and computer vision (CV), FMs [16] are general-purpose AI models that have superseded specialized, task-specific models across numerous applications. Following the recent success of FMs in earth sciences, we propose an FM trained on SDO data, which will be geared towards solving problems in solar and heliospheric physics. We believe that this is the first study to design a foundation model in the domain of heliophysics.

In Sec. 2, we discuss the current state-of-the art in ML applications in solving problems in solar physics and heliophysics. There we also discuss why FMs can provide better solutions to these problems, along with the specific downstream tasks we have planned to implement with the pre-trained FM. In Sec. 3, we describe an ML-ready dataset we have prepared to train the FM. In Sec. 4, we detail the FM architecture, along with strategies and methods used in the training. In Sec. 5, we show the results we have achieved during the early training process. We provide our discussion, conclusions, and future plans in Sec. 6.

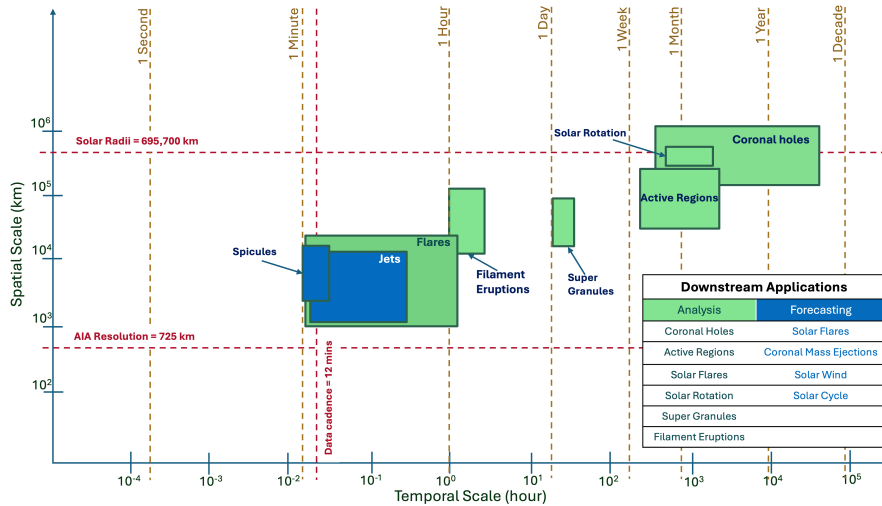


Fig. 2: Visualisation of selected downstream tasks with their spatial and temporal scales

2 Applications & State-of-the-Art Machine Learning Models

In recent years, the SDO data has been used by several ML-based methods to achieve results superior to conventional methods. The application of ML methods using SDO data is very broad in the field of solar and heliospheric physics. Here, we list some of the applications.

ML methods have been very successful in feature detection on the Sun, be it CHs, ARs, prominences, or magnetic flux elements. For example, U-Net architectures have been successfully used to detect CHs and ARs in AIA full-disk images [17, 18]. [19] used similar architectures to identify and track solar magnetic flux elements observed in HMI magnetograms. Convolutional Neural Networks (CNNs) have been used successfully to delineate CH boundaries using a combination of AIA and HMI images [20]. CNN-based architectures have also been used to detect prominences and sunspots, in addition to CHs [21]. Deep Neural Networks (DNNs) have demonstrated 99.9% accuracy in classifying solar phenomena such as quiet Sun, prominences, filaments, sunspots, and flare ribbons [22]. In addition, [23] has used an unsupervised ML method called pixel-wise k-means to detect CHs.

Forecasting solar flares is one of the areas heavily explored using ML methods trained on SDO data. Various studies have utilized non-potentiality proxies derived from LOS and vector magnetograms [24, 25], combining these with parameters from UV and EUV observations of ARs [26, 27]. Deep learning approaches have also been employed to forecast solar flares using magnetograms and intensitygrams without relying on external non-potentiality proxies of ARs [28]. Explainable models for flare prediction have been explored; for instance, [29] applied Gradient-weighted Class Activation Mapping (Grad-CAM) to identify influential regions on the solar surface

that contribute to flare predictions, which were found to be regions close to polarity inversion lines.

The SDO data has also been used by ML methods to predict solar wind and space weather at Earth. For example, solar wind speed prediction at Earth has been addressed by training DNNs on sequences of AIA images [30, 31]. In [32], CNN based deep learning model has been used for solar wind forecasting. Additionally, [33] developed a DNN using SDO/AIA 193 Å images to obtain probabilistic forecasts of geomagnetic activity, specifically the Kp index.

ML methods, especially deep learning techniques, in combination with the SDO data have been used for advancements in 3D reconstruction of the solar corona. For example, [34] employed a CNN to infer the Z -component of coronal loops solely from their two-dimensional shapes in EUV images. Meanwhile, [35] demonstrated that Generative Adversarial Networks (GANs) can map photospheric magnetograms to electron density distributions at various atmospheric heights.

Another category of applications that ML methods have made possible is the creation of synthetic data using image-image translation. For example, [36] trained a CNN with the SDO data to translate AIA images into disk-integrated irradiance data observed by SDO/EUV Variability Experiment (EVE; [37]). The CNN based architecture has also been used in autocalibration of SDO/AIA channels in [38]. In [39–41], Conditional GANs (cGANs) have been used to generate synthetic LOS magnetograms from Solar TERrestrial RELations Observatory (STEREO)/EUVI [42] images. The same cGANs were used to generate He-I images from SDO/AIA [43]. [44–46] explored the generation of synthetic EUV images in one wavelength from other wavelengths. In [47, 48] deep learning model has been used to translate SDO/HMI to SDO/AIA. Additionally, [49] proposed a CNN called *FarNet*, which associates solar far-side maps obtained through helioseismic holography with solar magnetograms obtained when the far-side has rotated to face the Earth.

ML-based methods have been used to improve the resolution of the existing data. For example, [50] introduced *Enhance*, a deep CNN that produces super-resolved continuum images and magnetograms for SDO/HMI, effectively increasing the per pixel resolution from $0.5''$ to $0.25''$. Similarly, [51] utilized GANs to improve the resolution of the Solar and Heliospheric Observatory (SOHO)/Michelson Doppler Imager (MDI; [52]) magnetograms, to match the resolution of SDO/HMI magnetograms.

Further, ML methods trained using SDO data have shown utility in solar cycle prediction [e.g. 53–57]. However, the results of these predictions vary widely and some of the predictions appear further from the truth than others, as the solar cycle 25 maxima approaches.

The applications listed above are not an exhaustive list of what is possible when ML methods are applied to SDO data. For a detailed review on this topic, see [58]. However, they convey that a broad range of solar and heliospheric problems can be targeted using the SDO data. In Fig. 2, we show the temporal and spatial scales of various solar structures and processes. We notice that the majority of these can be studied with an SDO dataset with full spatial resolution and a time cadence of 12 minutes. An FM model trained with such a database can be fine-tuned for various downstream applications. These include the analysis of solar features such as CHs,

ARs, and supergranules and processes such as solar flares, solar rotation, and filament eruptions. Besides analysis, an FM can be fine-tuned for forecasting specific downstream tasks like solar flare forecasting, filament eruption forecasting, and solar wind forecasting.

The FM approach envisioned in this work can be very beneficial to solve many of the outstanding problems in solar and heliospheric physics. Here, we list the potential benefits of using the FM approach:

1. **Mitigation of Supervision Bottleneck:** By leveraging pre-training on extensive solar datasets, FMs reduce the reliance on labeled data. This mitigates the supervision bottleneck and improves real-world performance, particularly for tasks with forecasting of rare events like large solar flares. The scarcity of the data creates a number of issues for robust prediction (and even the data-driven detection) of solar weather events using traditional methods. The introduction of FMs have a great potential to solve this problem by removing the hefty early training process and have diverse application areas.
2. **Versatility and Adaptability:** FMs can be easily fine-tuned for a variety of downstream tasks in solar and heliospheric physics, such as solar feature detection (coronal holes, sunspots, and active regions), space weather forecasting, and heliospheric modeling. The reduced need for labeled data and supervision makes these models adaptable to new datasets and tasks, providing significant value for researchers in the field.
3. **Improved Generalization:** These models exhibit superior generalization capabilities, enabling them to handle data distribution shifts. This is especially useful for predicting rare or extreme events, such as solar flares, filament eruptions, and CMEs, where traditional models may struggle with generalizing from training data.
4. **Scalability and Innovations:** FMs scale well with larger and higher-resolution datasets, making them ideal for handling multi-scale and multi-physics data from solar observations. This opens up new opportunities for high-resolution modeling of solar dynamics and heliospheric conditions, contributing to the next generation of data-driven models.
5. **Potential for Integration with Multi-Modal Data:** The emerging trend of FMs presents opportunities for integrating multi-modal data from various solar and heliospheric observatories, such as combining SDO, SOHO, Parker Solar Probe, and other ground-based or space-based observations. This integration promises to broaden the scope and accuracy of future predictions and analyses of solar activity, space weather events, and heliospheric conditions by leveraging diverse datasets that include magnetograms, EUV images, and in-situ measurements.

3 Data and Preparation

In this work, we train the FM using time series data comprising solar images captured in seven EUV wavelengths (304 Å, 171 Å, 193 Å, 211 Å, 335 Å, 94 Å, and 131 Å) by SDO/AIA, along with LOS and vector magnetograms obtained from SDO/HMI. Below, we describe how the SDO data was acquired, curated, and pre-processed to create an ML-ready dataset used for training the FM.

SDO data is publicly available through the Joint Science Operations Center (JSOC; <http://jsoc.stanford.edu>) as time series of various data products. The data products we have used to create our database are *aia.lev1_ewv_12s*, *hmi.M_720s*, and *hmi.B_720s*. The *aia.lev1_ewv_12s* series provides multi-wavelength AIA data with the resolution of 4096x4096 pixels and 12 second time cadence. The *hmi.M_720s* and *hmi.B_720s* series provide LOS and vector magnetograms, respectively, of the solar disk with the image resolution of 4096x4096 pixels and 12-minute cadence. To use HMI and AIA data together in the model, they should be aligned both spatially and temporally. The following subsections describe the necessary pre-processing required to achieve this:

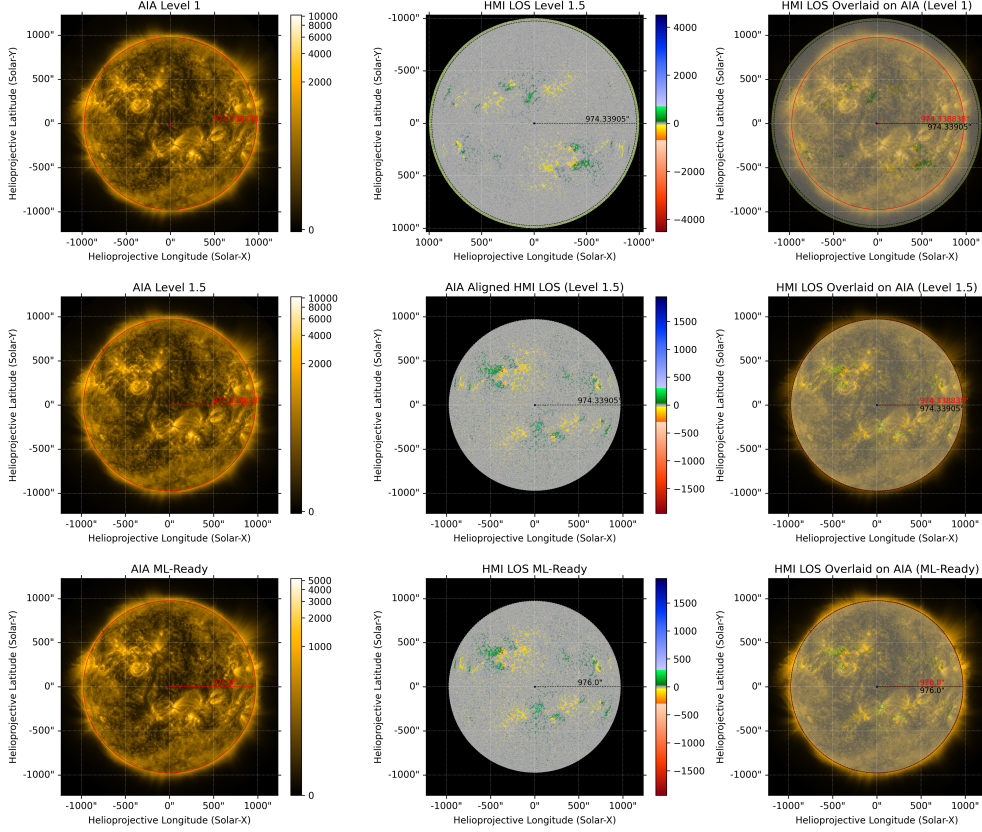


Fig. 3: Example of preprocessing steps of ML-Ready data on 2012-01-30 at 22:12 UT

3.1 AIA

The *aia.lev1_ewv_12s* series provided by JSOC contains level-1 data. This means that the images still include the roll angle of the satellite, i.e., the solar north-south axis is not aligned with the vertical y-axis, and each channel may have a slightly different

pixel scale. Hence it's important to promote the AIA data from level 1 to level 1.5 to analyze the images from all the bands together. The promotion to level-1.5 involves updating the pointing keywords, removing the roll angle, scaling the image to a resolution of 0.6 arcsec per pixel, and translating the image such that the center of the Sun is located in the center of the image. Besides these steps, exposure time normalization is an extra but necessary step during the promotion because AIA measurements have heterogeneous exposure time ranging from 0.05s to 2.9 seconds. We show the conversion of AIA data from level 1 to level 1.5 in the two top left panels of Figure 3 using an example of AIA 171 data.

Since the database contains data throughout the SDO lifetime, CCD camera sensor degradation also needs to be taken into account. The table of correction parameters calculated by the AIA science team is made publicly available via JSOC. These parameters are a time series of scalars that can be multiplied by the full disk AIA data to rectify the instrument degradation. However, this method is problematic when the corrected values in some of the pixels exceed the instrument saturation value (16383 for AIA). We post-process the degradation-corrected image to make sure that none of the pixels reach a value larger than this limit. In Figure 4, we plot the mean values of the full disk for each of the seven AIA wavelengths in level 1 (left panel) and level 1.5 with degradation correction (right panel) data. The variation of means over the years also reflects the solar cycle in which the high mean indicates strong solar activity. We notice that the degradation correction restores the higher activity in solar cycle 25 compared with the solar cycle 24.

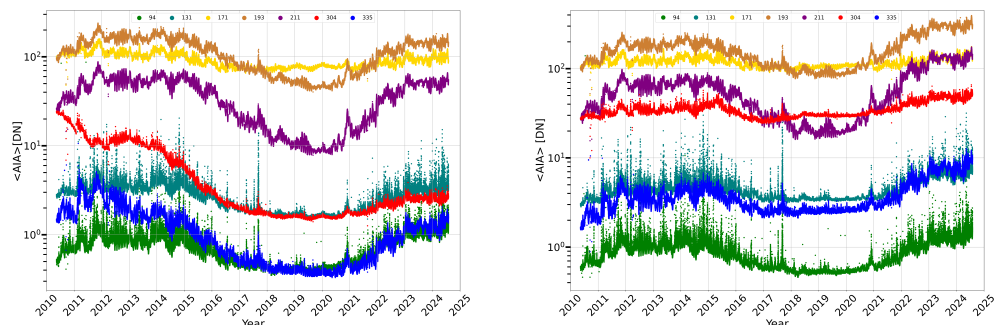


Fig. 4: Mean of full-disk AIA images per second in seven wavelengths over time for before (*left panel*) and after (*right panel*) the degradation correction.

One final step in making sure that the AIA data is ML-ready is to make the solar disk size the same in the whole database for all wavelengths, correcting for the elliptical orbit of the spacecraft. This step makes the solar disk radius as 976 arcsecs in all images. An example of this step is shown in the bottom panel of Figure 3.

3.2 HMI

Though the observable HMI data provided from JSOC are level 1.5, the data have a different orientation due to the way the HMI is mounted on SDO. Therefore the HMI data needs to be re-projected to be spatially aligned with the level 1.5 AIA images. This re-projection is done using the *reproject_exact* function available in SunPy [59]. This function preserves the magnetic flux when spatially aligning the HMI magnetograms with the AIA images. An example of the re-projection is shown in the top two panels in the middle column of Figure 3. Finally, similar to AIA images, the HMI images also needed to be corrected for elliptical orbit variation and fix the solar disk size to 976 arcsecs throughout the database. This step is shown with an example in the bottom panel in the middle column of Figure 3. The panels in the right column of this image show that after our processing, the solar disk in the AIA and HMI images are well aligned.

3.3 Temporal Alignment of HMI and AIA Data

In our database, we have decided to keep the temporal resolution of the time series at 12 minutes. This is because the low-noise LOS and vector magnetograms are available with this cadence from HMI. Moreover, as discussed earlier, and shown in Fig. 2, our targeted downstream tasks for the FM can be satisfactorily implemented using a database with this cadence. For each of the timestamps in *hmi.M_720s* series, we find the corresponding AIA data at that time and add it to our database. If the quality of the AIA data for any of the seven wavelengths is not good at an exact timestamp of HMI, we search for a time within 2 minutes of the timestamp that contains good quality AIA data in all seven EUV wavelengths and include that data in our database. If we are unable to find good-quality AIA data within the two minutes range, we do not include the data in our database. We consider the quality of the AIA data to be good if the QUALITY flag in the file header is equal to zero. Non-zero values indicate different reasons for invalid data. In Figure 5, we give examples of when the AIA data was found to be of bad quality.

3.4 Dataset Statistics

Our database contains ML-ready SDO data captured from May 13, 2010, to July 31, 2024. During this interval, there are about 2.9% data unavailable (18,261 out of 623,280 total timestamps) for *hmi.M_720s* series. The processed level 1.5 AIA and HMI data are stored in hourly netCDF files in float32 format, with data shape of [5, 8, 4096, 4096]. Each netCDF file is about 2.2 GB, and the total size of the data for training is about 257 TB.

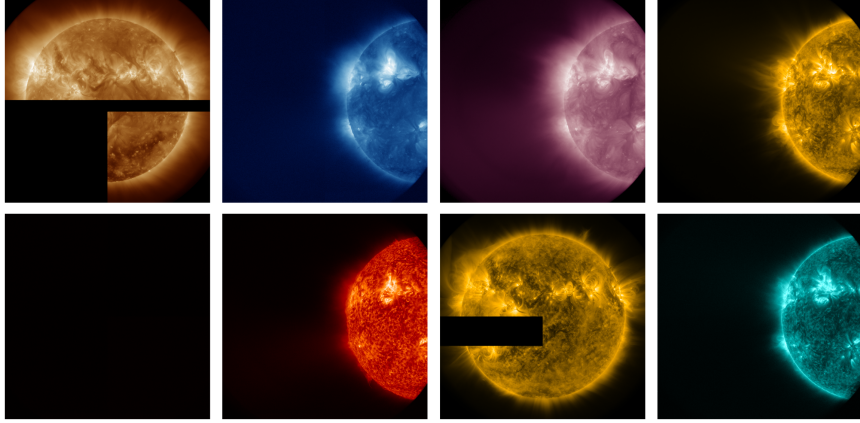


Fig. 5: Bad AIA Measurements due to a Variety of Reasons

For normalizing data of AIA and HMI we considered applying signum log over data given by:

$$\begin{aligned}
 f(x) &= \text{sgn}(x) \cdot \log(1 + |x|) \\
 f(x) &= \begin{cases} -\log(1 - x) & ; x \leq 0 \\ \log(1 + x) & ; x > 0 \end{cases} \\
 f^{-1}(x) &= \begin{cases} -e^{-x} + 1 & ; x \leq 0 \\ e^x - 1 & ; x > 0 \end{cases} \\
 f^{-1}(x) &= \text{sgn}(x) \left[e^{|x|} - 1 \right]
 \end{aligned}$$

4 Methodology

4.1 Sampling strategy

A key consideration when pretraining a FM is to use a representative dataset that contains sufficient examples of both common and rare events. This will enable the FM to efficiently learn characteristics and representations of all solar phenomena that could be relevant to downstream tasks. What constitutes common and rare events on the Sun depends on factors like where we are in the solar cycle and what the focus of a particular downstream task is. For example, solar minimum is characterized by periods of little solar activity, whereas solar maximum has ample activity in the form of ARs, CHs, flares, etc. [60]. However, events such as large CMEs are relatively rare in comparison with other events. The relevance of events, or regions on the Sun, also depends on the downstream application of interest, e.g. studies on magnetic flux would focus more on ARs and not consider granulation from quiet regions on the Sun [61], whereas studies of the solar wind would focus on various explosive events, but in particular on CHs [62]. This highlights the need for a dataset that isn't biased towards any prevalent phenomenon and that accounts for spatial and temporal dynamics.

Despite ongoing efforts, determining the optimal composition of training data in computer vision remains elusive. Most current practices adjust the data mixture heuristically, often increasing the proportion of high-quality images or underrepresented classes without providing detailed criteria [14]. Additionally, it’s challenging to know whether these data strategies will be effective before completing the training process.

We consider the task of pretraining a Vision Transformer (ViT) model p_θ for image modeling tasks such as masked image reconstruction or autoregressive prediction. The training dataset is defined as $\mathcal{D}_{\text{train}} = \{\mathcal{D}_i\}_{i=1}^M$, comprising M visual domains. The training dataset contains Z events (images), with each domain \mathcal{D}_i containing Z_i events such that:

$$Z = \sum_{i=1}^M Z_i \quad (1)$$

We define z as the total number of events to be sampled from the dataset $\mathcal{D}_{\text{train}}$ during training. Each domain \mathcal{D}_i will contribute z_i events to the sampled set, where:

$$z = \sum_{i=1}^M z_i \quad (2)$$

The mixture proportions are defined as:

$$r_i = \frac{z_i}{z} \quad (3)$$

These proportions determine the probability of sampling an image from domain \mathcal{D}_i . We sample images \mathbf{x} from the entire dataset $\mathcal{D}_{\text{train}}$ based on the mixture proportions $r = [r_1, r_2, \dots, r_M]$. Each image \mathbf{x} is selected from domain \mathcal{D}_i with probability r_i .

The Vision Transformer (ViT) model processes each image as a sequence of these patches. The loss function incorporates the sampling approach based on z events and is defined over the sampled images. For masked image modeling where a subset of patches \mathcal{M} is masked, the model aims to reconstruct the original pixel values of these masked patches.

$$\mathcal{L}_\theta = \frac{1}{z} \sum_{i=1}^M \sum_{\mathbf{x} \in \mathcal{D}_i} \mathbb{I}_{\text{sampled}}(\mathbf{x}) \left[- \sum_{j \in \mathcal{M}} \log P_\theta(x_j | \mathbf{x}_{\setminus \mathcal{M}}) \right] \quad (4)$$

where \mathcal{M} is the set of indices for the masked patches in image \mathbf{x} , x_j is the pixel content of the masked patch at position j , $\mathbf{x}_{\setminus \mathcal{M}}$ represents the image with patches in \mathcal{M} masked out, $P_\theta(x_j | \mathbf{x}_{\setminus \mathcal{M}})$ is the model’s probability of reconstructing the original patch x_j given the unmasked portions of the image and the loss is summed over the masked patches and averaged over z .

In case of Autoregressive Modeling,

$$\mathcal{L}_\theta = \frac{1}{z} \sum_{i=1}^M \sum_{\mathbf{x} \in \mathcal{D}_i} \mathbb{I}_{\text{sampled}}(\mathbf{x}) \left[- \sum_{j=1}^N \log P_\theta(x_j | x_{1, \dots, j-1}) \right] \quad (5)$$

where $N = 65,536$ is the total number of patches per image, x_j is the j -th patch in the sequence, x_1, \dots, x_{j-1} are all preceding patches, $P_\theta(x_j | x_1, \dots, x_{j-1})$ is the probability of predicting patch x_j given all previous patches and the loss is summed over all patches in the image and averaged over z .

4.2 Training objectives

Foundation models are trained with task-agnostic pretext tasks. That is, training tasks that are intrinsic to the data and do not depend on any downstream use case yet that let the model learn a representation that can be applied to a large number of downstream tasks. For vision models, the canonical pretext tasks are masked reconstruction [63] as well as contrastive learning [64]. In the former case, the model is given an input image with large parts missing and has then to reconstruct the entire image. The approach is easily implemented on transformer architectures. While generalization to higher dimensional settings such as videos is straightforward [65], this comes at a cost in token counts which can become problematic for standard ViT architectures. Contrastive approaches typically depend on randomized augmentations of the data. The pretext task is then to align embeddings of data irrespective of augmentation. This has been found to be particularly useful in the multi-modal setting.

Masked reconstruction and contrastive learning are widely used pretext tasks. While this is partially due to the powerful representations one obtains using them, another reason for their wide adoption is also that they easily apply to any image dataset. Yet in cases where the underlying dataset has more structure than just that of an image or data cube it is possible to consider domain specific pretext tasks. Ideally these should mirror some properties of the domain to prompt the learning of more powerful and domain-specific representations. A nice example of this is the pretext task of AtmoDist [66]. Here, the authors considered weather data from the ERA5 reanalysis. Their model was pretrained to classify the temporal difference between data from two different time stamps. If we stick to the weather and climate domain, [15] was pretrained to predict data at future or present times given incomplete knowledge about the present. This combines masking with forecasting albeit in a different way to the work by [65] cited above. Moreover, it reflects some fundamental properties of weather data: That data is naturally sparse as it is the output of a complex observational system consisting of weather stations, satellites etc. Yet at the same time the forecasting problem and the associated understanding of temporal dynamics is absolutely fundamental as well.

Beyond weather and climate data one can look at satellite imagery. The Prithvi model of [14] was trained with spatial masking. However, it is straightforward to generalize the masking principle to the reconstruction of certain satellite bands.

With this in mind we can finally turn to solar data. Clearly dynamics play an important role so one would expect forecasting to be a powerful pretext task. At the same there is a strong interest in instrument emulation as shown by the work of [67]. Which of course raises the possibility of using instrument to instrument translation as a non-temporal pretext task. Naturally one should also consider masking and contrastive learning. Masking is particularly useful if one aims for a model that can operate on different spatial domains. That is, if the domain of the data at inference time might

change from the whole disk to a subset of the sun, pretraining with a masking objective ensures that the model is “comfortable” with missing data. This might be of particular interest in the context of heliophysics where instruments generally observe the part of the sun accessible via line of sight. All in all one expects that some combination of masking, forecasting and instrument to instrument translation can be used as a powerful pretext task to ensure that a foundation model for the sun learns a powerful representation of instrument data that can be applied to many use cases.

4.3 Vision Transformers

Recent studies on vision transformer-based architectures have introduced various strategies to implement attention mechanisms. Attention layers contribute to model accuracy by capturing long-range dependencies, handling diverse data types, accommodating various biases, and allowing for parallel processing. The strategies adopted in recent works for attention implementation can be broadly divided into two categories, namely, 3D self-attention and 2D cross-attention. The key advantage of these approaches is their focus on reducing the heavy computational demands typically associated with attention mechanisms in this data-rich field. Techniques like window-based attention, as employed in the Swin transformer [68], reduce the computational load by dividing feature maps into windows and applying self-attention within each window. This method has proven effective in both non-spherical and spherical scenarios, showing strong performance. Another approach leverages variable aggregation and cross-attention to simplify the attention complexity found in the basic Vision Transformer (ViT) architecture [69]. We utilized these attention-based strategies to train our model. In next few paragraphs we will discuss details about the approaches taken by us and corresponding outcomes.

Before delving into the details of these attention strategies followed by us, one must note that one of the approaches that has been utilized majorly for pretraining a foundation model is masking. In this paradigm, a certain portion of the input data is masked, and the model learns to recreate the masked image. Moreover, masking is done in terms of patches of size $H \times W$, where H and W are dimensions in pixels. The reason behind the use of masking is firstly because it lowers the memory load of the model on the GPU so that one can fit larger batch sizes and secondly, it helps the model to learn spatial information effectively. However, token size becomes a constraint eventually for such a large model. Even large vision transformer models like Prithvi WxC [15], which has around 2.3 billion parameters, uses approximately 60,000 tokens (the input image dimension is 576x360 and each token size is 2x2). Similarly, if we try to do the same with SDO image with 8 or 10 channels with 4K resolution - we will end up with more than 8 million tokens. However, training such a large model directly is very difficult as per current hardware configurations. To counter these challenges and improve local and global attention [70] et al. came up with a shifted window mechanism, where we pay attention to tokens specified in a window.

Based on this logic we tried to train an architecture *Model1* inspired by [71] where we masked 50% of the image. There were two objectives of the model- one to predict the mask and the second to generate an image given a lead time. Lead time is defined as ΔT where T can be between 1-5, i.e. 12-60 mins. Figure 6 shows the architecture used

to train the model is heavily inspired from [70]. The number of trainable parameters in this model is 1.8 billion, where the patch size is $16 \times 16 \times 2$. The model is given input of 2 timesteps. However, the experiments results were not good at all, with 30 epochs of training on 4 years of data. We also tried different masking strategy ranging from 10% to 80% at interval of 10.

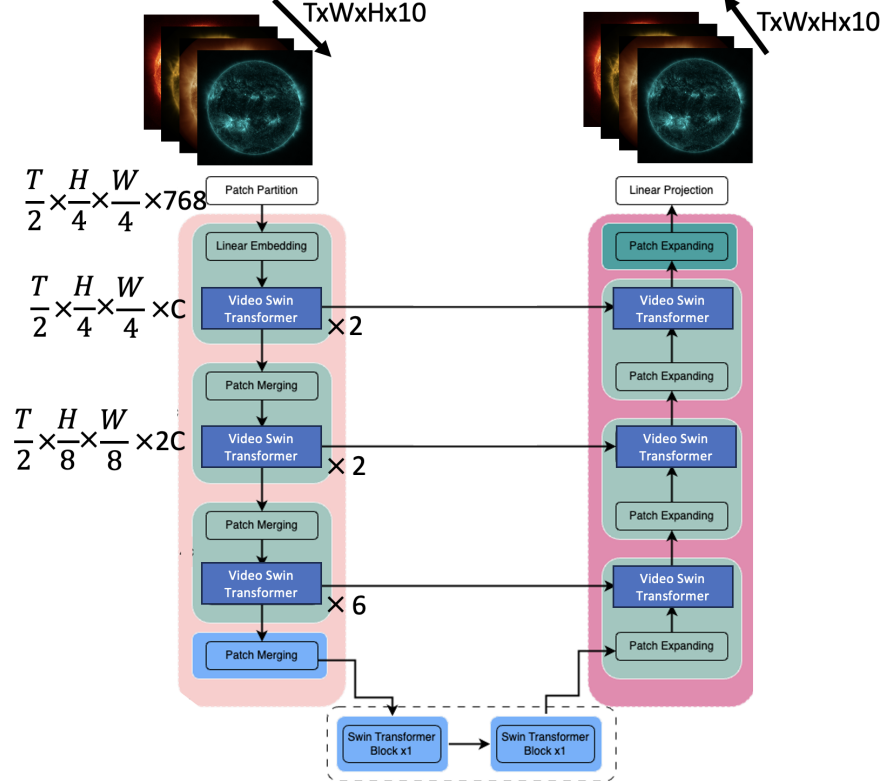


Fig. 6: Architecture of *Model1*, which uses video swin transformer blocks with U-Net alike structure for training on SDO data. The model takes two input time steps and forecast the next time step.

However, the model trained with shifted windowed attention fails at generating the mask patches quickly. Hence, we tried another approach inspired by [72]. In this model, token mixing strategy is used by spectral blocks. This approach of token mixing makes the model faster. Using this methodology, post the patch embedding stage, the self-attention mechanism is replaced with a non-parametric Fourier transform, followed by a non-linear activation and a feed-forward network. This approach shows faster convergence to the model and better results for prediction for early epochs of training. The object of this model is to generate a lead time step only, i.e. forecasting model for time t . However, in this approach, masking technique can not be used as

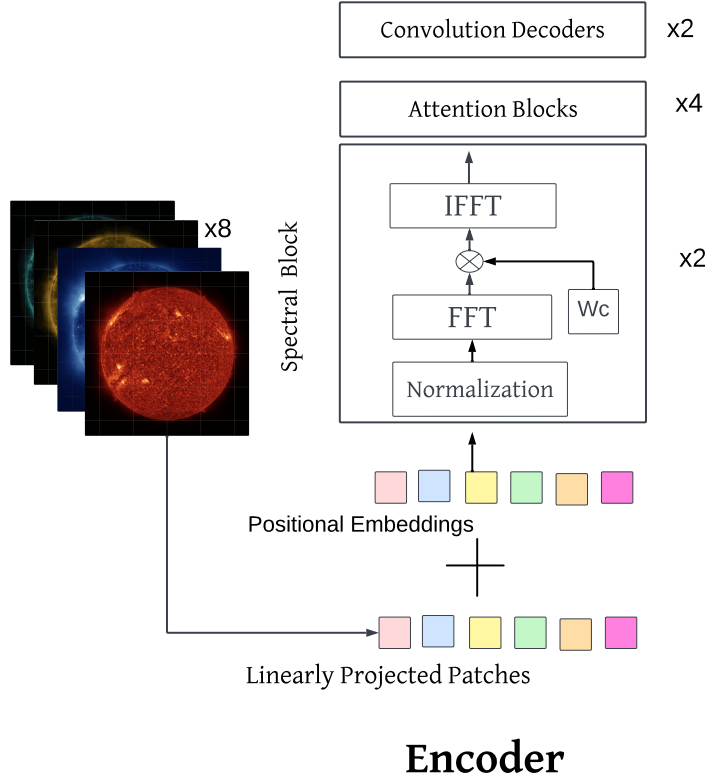


Fig. 7: Architecture of spectral transformer as described in [72] used for training on SDO data. Convolution Blocks are added as decoders for the next step generation.

there will be missing information for Fast-Fourier Transforms (FFT). Figure 7 shows the basic architecture design of this model. Though the authors proposed an encoder-only model, a couple of convolution blocks are attached by us at the end to form decoder. We kept the convolution decoder shallow to enforce maximum learning for representation by encoders. In this case, the model is trained on 64×64 token size and the embedding space is of 1024. However, it is previously mentioned that number of token going as low as 2×2 results in around 8 million tokens. To counter the issue, we adopted a hybrid architecture, explained in next section.

4.3.1 Long-Short Spectral Transformer

To have larger number of sequence tokens, we came up with another strategy of mixing tokens and removing regular attention with the approach proposed by [73]. In this approach, dynamic projection-based long-range attention is combined with

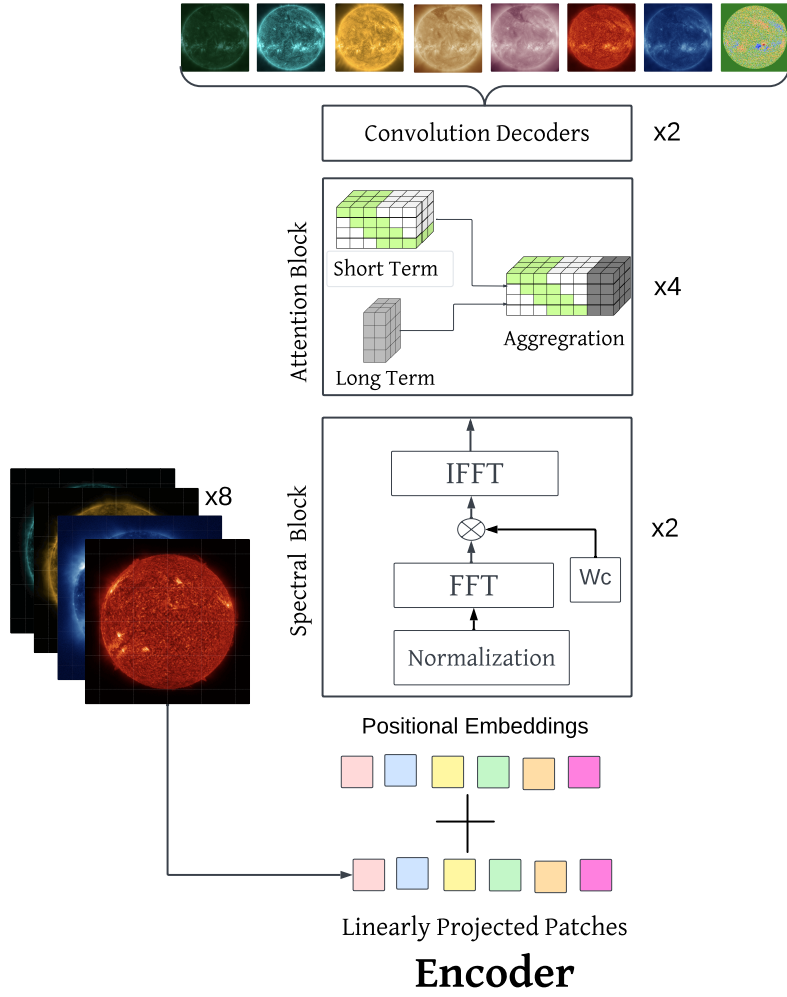


Fig. 8: Architecture of Long-Short Spectral Transformer, where we are using a learnable weight parameter (W_c) after performing FFT for estimating weights of frequency components. Post that we do inverse FFT to bring the information to physical space. The attention head is designed on the principle of long-short attention, where we calculate short-term attention by sliding window, and long-range attention by dynamic projections

local windowed short-term attention to effectively capture both global and local features in the input data. Specifically, long-range attention reduces the dimensionality of the Key and Value embeddings by projecting them into a smaller set of tokens using a projection matrix derived from the original Key embeddings. This mechanism allows the model to efficiently manage distant correlations across the sequence. On

the other hand, short-term attention divides the input sequences into non-overlapping segments, focusing on capturing fine-grained local correlations. By merging these two mechanisms—long-range dynamic projection for distant dependencies and short-term attention for detailed local interactions—the model successfully addresses different scales of feature representation. By utilizing this approach we were able to train the model with a 16×16 patch size with a batch size of 4 on A100 80 GB GPUs. Figure 8 shows the modified architecture of the model. In future, we will try to scale it to $8x8$ token size.

4.4 Diffusion Models

All the models discussed so far are largely deterministic in nature, i.e., given a pretext task such as masking, forecasting or some combination thereof one can train a transformer model on an RMSE objective or equivalent. While this approach is powerful and has led to many excellent results, it does have a few, known shortcomings.

To start, it is known that forecast models trained on an RMSE objective exhibit *blurring* at longer lead times. That is, if one looks reasonably far into the future individual features start to lose focus. This should not be surprising. Intuitively speaking one can think of the space of all possible futures given an initial state. If this space of small futures depends on small changes in initial conditions, the model will have to effectively “hedge its bets” when making a prediction. It does so by predicting some averaged state between possible futures. Or to say this differently: In light of an uncertain future the optimal deterministic model trained on an RMSE (MAE) objective yields the average (median) of all possible futures.

A related issue is the “double penalty” of the RMSE objective. That is, if there are small, well defined features, a model trained on RMSE is heavily penalized if it misplaces such features by a few pixels: It is penalized both at the true location of the feature where it underpredicts as well as at the falsely predicted location where it overpredicts. Thus, in light of extreme values and uncertainty a model is incentivized to predict a smaller value over a larger area.

Both these issues can be addressed by moving to diffusion-based generative models [74]. Roughly speaking, diffusion models train a denoising process that is capable of turning random noise into data samples. This denoising process can be conditioned on auxiliary data which allows us to model conditional distributions $p(x|y)$. The authors of [75, 76] used this to train probabilistic forecast models $p(x_t|x_{t-1})$. Early applications for SDO data condition on the intensity of solar flares with the aim of generating synthetic solar flares [77].

As the results of [75] show, in particular when compared with [78], such model are much more capable of modeling fine structure and extremes in the data. In particular see [79] showing a diffusion model trained at 1 degree resolution exhibiting better fine structure than a deterministic model trained at 0.25 degrees resolution.

Diffusion models have some additional salient features. The authors of [80] showed that diffusion models have implicit knowledge of the normal vector to the data manifold. This is of particular interest in the context of the question whether data-driven models accurately reflect the dynamics – the physics – of the underlying data. While

conditioning on physically consistent data – such as prior time steps or the measurements of other sensors – massively reduces the risk of generating unphysical data, it is reassuring to know that such models have for the lack of a better word an aversion to generating samples away from the (training) data distribution. On a related note one should consider the work of [81]. Here the authors used a diffusion process to improve the stability of neural PDE solvers during long rollouts. Essentially the diffusion model was trained to remove unphysical frequencies from AI-generated solutions to PDEs.

4.5 Graph methods

So far we have tacitly assumed that the data we are dealing with comes in the form of two-dimensional images or sequences thereof. This is a decent assumption since SDO data can easily be understood in image form and one can leverage results from the mainstream of computer vision architecture as well as the associated technological innovations and accelerations in deep learning toolkits. At the same time it is clear that the understanding of solar data in terms of images is limited. Large parts of the image model empty background space which is a waste of compute capacity. Moreover the structure of two-dimensional cartesian grids really does not respect the known structural priors of the data – that we are modeling a spherical body for which we have data from line-of-sight instruments.

While transformers technically do not see gridded images but unordered sets of tokens, the known scaling issues of attention mechanisms used in transformers mean that such architectures often use optimizations that do tacitly assume that data lives on a rectangular grid. (See however [15].) One alternative is to move to graph-based models. Here it is possible to define custom meshes and project to project the data on. This was done with great success in the work of [78]. In this paper the authors built a highly performant AI emulator for weather forecasting systems. At the heart of the emulator is an earth-encircling icosahedral hierarchical mesh. A graph neural network (GNN) then maps the data from a rectangular lat-lon grid onto this mesh, successively processes it through multiple layers and finally projects the data back onto the original grid.

While it is fairly straightforward to consider this application for solar physics – either by projecting partial observations from instruments onto a spherical mesh or by having a 2d disc with changing data density from the center to the edges – one should note that there are drawbacks. To start, the definition of such a mesh and network can be a considerable technical challenge. Moreover, GNNs while a very active research topic are not as much in the mainstream of AI as transformer architectures. Thus issues as their scaling performance are less understood.

5 Results

Initial results with Long-Short Spectral Transformer

Design of experiment: To calculate initial results and the capability of the model to learn quickly, we designed experiments that would run for a maximum of 24-36 hours with a maximum of 16 nodes of GPU resources where each node had 4 A100

80GB GPU. For the results below the model was trained on 4 years of SDO data (2011-2015 and validation was performed on 6 months of data. Data was normalized by using signum log as described in section 3.4.

Figure 9 shows the ground truth, prediction, structural similarity index measures (SSIM) map, and the squared difference of both by using our modified long-short Spectral transformer 8. The squared difference between the image is defined as $X = (Y - \hat{Y})^2$. The SSIM value for band 0171 is 0.83, for band 0193 is 0.90, and for band 0211 is 0.86. For other bands the SSIM was between 0.4 to 0.65. The root mean squared error (RMSE) for band 0171 is 0.11, for band 0193 is 0.095, and for band 0211 is 0.10. With just 20 epochs of training on 4 years of data, the model shows remarkable performance. It is to be noted that, further training will be able to refine the image structure and also remove very light patchy artifacts from the model, which is primarily influenced by token size.




Fig. 9: Forecast of AIA channels for next time steps, using 2 input timesteps. The first row is the ground truth for the model. The second row is the prediction from the model. The third row is the SSIM map between the ground truth and the prediction image. The fourth row is the squared error between the ground truth and the prediction image.

Figure 10 shows the prediction of the HMI channel by using the same model, where the first image from left is the prediction, followed by ground truth, squared error, and SSIM map. To look at a finer scale, we performed a center crop of 1K resolution on the HMI band and performed the same comparison as shown in figure 11 and the

images follow the same order as in 10. The SSIM score for the predicted image with respect to ground truth is 0.73 and the RMSE is 0.65.




Fig. 10: Forecast of HMI channels for next time steps. Ground truth denotes the image at time step t in SDO data and prediction is the corresponding prediction from the model. [From left] The first image is the ground truth, second is the prediction by the model. The third image is the squared difference between ground truth and prediction by the model and the fourth image is the SSIM map between ground truth and prediction.




Fig. 11: [From Left] Center crop of 1K resolution of (first image) ground truth, (second) prediction by the model. (third) squared difference between ground truth and prediction by the model and (fourth) SSIM map between ground truth and prediction.

Initial Results with Diffusion




Fig. 12: Generation of AIA-EUV channels through Latent Diffusion conditioned on HMI. Top row: 3-channel AIA-EUV image (channels 094, 131, and 171 displayed as RGB), corresponding HMI image used for conditioning the latent diffusion model, denoised sample generated from Gaussian noise, and reconstructed autoencoder AIA-EUV image. AIA-EUV and reconstructed images are shown for comparison only (other instances of these are used during training). Bottom row: Denoising sequence (7 timesteps) showing how the sample emerges from noise through successive denoising operations (1000 timesteps) conditioned on HMI.

In our current approach to diffusion models, we follow R. Rombach et al. [82] and first train shallow autoencoders to reduce image sizes before they are fed to the diffusion model. This latent diffusion model approach has been shown to “reach a near-optimal point between complexity reduction and detail preservation”, [82] and will allow us to eventually scale the model to the native 4096×4096 pixel resolution of SDO data. Initial experiments are done on downsampled versions and/or random crops of 3-channel SDO data (AIA-EUV channels 094, 131, and 171).

Among the many possibilities for conditioning the latent diffusion model, such as through class labels (e.g. strength of flares), semantic maps (e.g. position of active regions), or image-based conditioning (e.g. prior SDO timesteps, different SDO channels or different instruments), we chose one of the more challenging tasks: translation from HMI to AIA-EUV. HMI images are auto-encoded like the 3-channel SDO data or simply downsampled to the latent space dimensions through bilinear interpolation. We tried conditioning using either cross-attention or concatenation. Better results were achieved using concatenation.

Fig. 12 shows an example of AIA-EUV generation conditioned on HMI. Here we used a 4-level autoencoder to encode AIA-EUV images of $512 \times 512 \times 3$ pixels at $32 \times 32 \times 3$ pixels in latent space, where the latent diffusion model is applied. The HMI images were downsampled through bilinear interpolation to 32×32 pixels. Conditioning on HMI was achieved through concatenation.

6 Discussion & Conclusion

We presented the first study to design a foundation model in the domain of heliophysics. The development of a foundation model (FM) for heliophysics represents a significant step in leveraging deep learning techniques to advance our understanding of solar phenomena. One of the significant advantages of using a foundation model in heliophysics is the mitigation of the supervision bottleneck. Traditional supervised learning methods require extensive labeled datasets, which are challenging to obtain in heliophysics due to the rarity and complexity of certain solar events. By employing self-supervised learning during pre-training, the FM learns intrinsic features from the vast unlabeled SDO dataset, reducing the dependency on labeled data for downstream applications. This approach not only enhances the model’s adaptability but also allows for better generalization across different types of solar phenomena. In this paper, we have tried to draw a scope for designing a foundation model using the SDO dataset. We have also shown preliminary experimental results by pre-training an encoder-decoder architecture on four years of high-resolution data from the Solar Dynamics Observatory (SDO). We have aimed to capture the intricate temporal and spatial patterns inherent in solar activities. The inclusion of eight bands — seven EUV channels from the Atmospheric Imaging Assembly (AIA) and one magnetogram from the Helioseismic and Magnetic Imager (HMI) — provides a comprehensive dataset that encapsulates various information of the solar atmosphere.

Considering state-of-the-art methods like Swin transformer [68], Vision Transformer [69], and spectral transformer [72] and the challenges associated with them, we proposed a new approach for model design. Keeping note of image size and data volume and limitations for larger token size and lower dimensions, we also presented our approach that combines frequency-based information with long-short attention mechanisms, enabling the model to handle the dynamic nature of solar events. The integration of frequency-based components allows the model to capture local features, as well as global features for long-range dependencies. Additionally, by adding a long-short attention mechanism, we intend to maximize the visibility of more tokens by the model.

Despite the experimental constraints to find an optimized model, initial results are promising, demonstrating the model’s capability to predict solar dynamics and its potential applicability to various downstream tasks. For instance, the FM shows proficiency in capturing the evolution of active regions and predicting solar flares by analyzing patterns in magnetic field configurations and coronal activities. This ability is critical for space weather forecasting, where accurate predictions of solar events can mitigate the impacts on satellite operations, communication systems, and power grids on Earth.

In conclusion, the development of an AI foundation model for heliophysics holds significant promise for advancing solar research and improving space weather forecasting. While challenges related to data volume, computational resources, and variability exist, the initial results are encouraging. By addressing these challenges through efficient data handling, robust training strategies, and comprehensive validation, the FM can become a versatile tool for the heliophysics community. Ongoing efforts will focus on refining the model, expanding its applications, and ultimately enhancing our ability to predict and understand complex solar phenomena that impact our technological society.

Acknowledgements. The Authors acknowledge the National Artificial Intelligence Research Resource (NAIRR) Pilot and NVIDIA for providing support under grant no. NAIRR240178. The authors would also like to thank Jülich Supercomputing Centre (JSC) and NASA Advanced Supercomputing (NAS) Division for their compute support.

References

- [1] Pesnell, W.D., Thompson, B.J., Chamberlin, P.C.: The Solar Dynamics Observatory (SDO). *Solar Physics* **275**(1-2), 3–15 (2012) <https://doi.org/10.1007/s11207-011-9841-3>
- [2] SDO Publications (2024). <https://sdo.gsfc.nasa.gov/mission/publications.php>
- [3] Schou, J., Scherrer, P.H., Bush, R.I., Wachter, R., Couvidat, S., Rabello-Soares, M.C., Bogart, R.S., Hoeksema, J.T., Liu, Y., Duvall, T.L., Akin, D.J., Allard, B.A., Miles, J.W., Rairden, R., Shine, R.A., Tarbell, T.D., Title, A.M., Wolfson, C.J., Elmore, D.F., Norton, A.A., Tomczyk, S.: Design and Ground Calibration of the Helioseismic and Magnetic Imager (HMI) Instrument on the Solar Dynamics Observatory (SDO). *Solar Physics* **275**(1-2), 229–259 (2012) <https://doi.org/10.1007/s11207-011-9842-2>
- [4] Lemen, J.R., Title, A.M., Akin, D.J., Boerner, P.F., Chou, C., Drake, J.F., Duncan, D.W., Edwards, C.G., Friedlaender, F.M., Heyman, G.F., Hurlburt, N.E., Katz, N.L., Kushner, G.D., Levay, M., Lindgren, R.W., Mathur, D.P., McFeaters, E.L., Mitchell, S., Rehse, R.A., Schrijver, C.J., Springer, L.A., Stern, R.A., Tarbell, T.D., Wuelser, J.-P., Wolfson, C.J., Yanari, C., Bookbinder, J.A., Cheimets, P.N., Caldwell, D., Deluca, E.E., Gates, R., Golub, L., Park, S., Podgorski, W.A., Bush, R.I., Scherrer, P.H., Gumm, M.A., Smith, P., Auken, G., Jerram, P., Pool, P., Souffi, R., Windt, D.L., Beardsley, S., Clapp, M., Lang, J., Waltham, N.: The Atmospheric Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO). *Solar Physics* **275**(1-2), 17–40 (2012) <https://doi.org/10.1007/s11207-011-9776-8>
- [5] Heinemann, S.G., Temmer, M., Heinemann, N., Dissauer, K., Samara, E., Jerčić, V., Hofmeister, S.J., Veronig, A.M.: Statistical Analysis and Catalog of Non-polar Coronal Holes Covering the SDO-Era Using CATCH. *Solar Physics* **294**(10),

- 144 (2019) <https://doi.org/10.1007/s11207-019-1539-y> arXiv:1907.01990 [astro-ph.SR]
- [6] van Driel-Gesztelyi, L., Green, L.M.: Evolution of Active Regions. *Living Reviews in Solar Physics* **12**(1), 1 (2015) <https://doi.org/10.1007/lrsp-2015-1>
- [7] Reale, F.: Coronal loops: Observations and modeling of confined plasma. *Living Reviews in Solar Physics* **11**(1), 4 (2014) <https://doi.org/10.12942/lrsp-2014-4>
- [8] Liu, W., Ofman, L.: Advances in Observing Various Coronal EUV Waves in the SDO Era and Their Seismological Applications (Invited Review). *Solar Physics* **289**(9), 3233–3277 (2014) <https://doi.org/10.1007/s11207-014-0528-4> arXiv:1404.0670 [astro-ph.SR]
- [9] Benz, A.O.: Flare observations. *Living Reviews in Solar Physics* **14**(1), 2 (2016) <https://doi.org/10.1007/s41116-016-0004-3>
- [10] Parenti, S.: Solar prominences: Observations. *Living Reviews in Solar Physics* **11**(1), 1 (2014) <https://doi.org/10.12942/lrsp-2014-1>
- [11] Mukkavilli, S.K., Civitarese, D.S., Schmude, J., Jakubik, J., Jones, A., Nguyen, N., Phillips, C., Roy, S., Singh, S., Watson, C., et al.: Ai foundation models for weather and climate: Applications, design, and implementation. arXiv preprint arXiv:2309.10808 (2023)
- [12] Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J.K., Grover, A.: ClimaX: A foundation model for weather and climate. arXiv (2023) [2301.10343](https://arxiv.org/abs/2301.10343)
- [13] Bodnar, C., Bruinsma, W.P., Lucic, A., Stanley, M., Brandstetter, J., Garvan, P., Riechert, M., Weyn, J., Dong, H., Vaughan, A., et al.: Aurora: A foundation model of the atmosphere. arXiv preprint arXiv:2405.13063 (2024)
- [14] Jakubik, J., Roy, S., Phillips, C.E., Fraccaro, P., Godwin, D., Zadrozny, B., Szwarcman, D., Gomes, C., Nyirjesy, G., Edwards, B., Kimura, D., Simumba, N., Chu, L., Mukkavilli, S.K., Lambhate, D., Das, K., Bangalore, R., Oliveira, D., Muszynski, M., Ankur, K., Ramasubramanian, M., Gurung, I., Khallaghi, S., Hanxi, Li, Cecil, M., Ahmadi, M., Kordi, F., Alemohammad, H., Maskey, M., Ganti, R., Weldemariam, K., Ramachandran, R.: Foundation Models for Generalist Geospatial Artificial Intelligence (2023)
- [15] Schmude, J., Roy, S., Trojak, W., Jakubik, J., Civitarese, D.S., Singh, S., Kuehnert, J., Ankur, K., Gupta, A., Phillips, C.E., Kienzler, R., Szwarcman, D., Gaur, V., Shinde, R., Lal, R., Silva, A.D., Diaz, J.L.G., Jones, A., Pfreundschuh, S., Lin, A., Sheshadri, A., Nair, U., Anantharaj, V., Hamann, H., Watson, C., Maskey, M., Lee, T.J., Moreno, J.B., Ramachandran, R.: Prithvi WxC: Foundation Model for Weather and Climate (2024). <https://arxiv.org/abs/2409.13598>

- [16] Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., et al.: On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258 (2021)
- [17] Illarionov, E.A., Tlatov, A.G.: Segmentation of coronal holes in solar disc images with a convolutional neural network. *Monthly Notices of the Royal Astronomical Society* **481**(4), 5014–5021 (2018) <https://doi.org/10.1093/mnras/sty2628> <https://academic.oup.com/mnras/article-pdf/481/4/5014/25992150/sty2628.pdf>
- [18] Mackovjak, Š., Harman, M., Maslej-Krešňáková, V., Butka, P.: Scss-net: solar corona structures segmentation by deep learning. *Monthly Notices of the Royal Astronomical Society* **508**(3), 3111–3124 (2021)
- [19] Jiang, H., Wang, J., Liu, C., Jing, J., Liu, H., Wang, J.T.L., Wang, H.: Identifying and tracking solar magnetic flux elements with deep learning. *The Astrophysical Journal Supplement Series* **250**(1), 5 (2020) <https://doi.org/10.3847/1538-4365/aba4aa>
- [20] Jarolim, R., Veronig, A., Hofmeister, S., Heinemann, S., Temmer, M., Podladchikova, T., Dissauer, K.: Multi-Channel Coronal Hole Detection with Convolutional Neural Networks. In: EGU General Assembly Conference Abstracts. EGU General Assembly Conference Abstracts, pp. 21–1490 (2021). <https://doi.org/10.5194/egusphere-egu21-1490>
- [21] Baek, J.-H., Kim, S., Choi, S., Park, J., Kim, J., Jo, W., Kim, D.: Solar event detection using deep-learning-based object detection methods. *Solar Physics* **296**(11), 160 (2021)
- [22] Armstrong, J.A., Fletcher, L.: Fast solar image classification using deep learning and its importance for automation in solar physics. *Solar Physics* **294**(6), 80 (2019) <https://doi.org/10.1007/s11207-019-1473-z>
- [23] Inceoglu, F., Shprits, Y.Y., Heinemann, S.G., Bianco, S.: Identification of coronal holes on aia/sdo images using unsupervised machine learning. *The Astrophysical Journal* **930**(2), 118 (2022)
- [24] Ahmed, O.W., Qahwaji, R., Colak, T., Higgins, P.A., Gallagher, P.T., Bloomfield, D.S.: Solar flare prediction using advanced feature extraction, machine learning, and feature selection. *Solar Physics* **283**(1), 157–175 (2013) <https://doi.org/10.1007/s11207-011-9896-1>
- [25] Bobra, M.G., Couvidat, S.: SOLAR FLARE PREDICTION USING SDO/HMI VECTOR MAGNETIC FIELD DATA WITH a MACHINE-LEARNING ALGORITHM. *The Astrophysical Journal* **798**(2), 135 (2015) <https://doi.org/10.1088/0004-637x/798/2/135>

- [26] Nishizuka, N., Sugiura, K., Kubo, Y., Den, M., Watari, S., Ishii, M.: Solar flare prediction model with three machine-learning algorithms using ultraviolet brightening and vector magnetograms. *The Astrophysical Journal* **835**(2), 156 (2017) <https://doi.org/10.3847/1538-4357/835/2/156>
- [27] Nishizuka, N., Sugiura, K., Kubo, Y., Den, M., Ishii, M.: Deep flare net (DeFN) model for solar flare prediction. *The Astrophysical Journal* **858**(2), 113 (2018) <https://doi.org/10.3847/1538-4357/aab9a7>
- [28] Abed, A.K., Qahwaji, R., Abed, A.: The automated prediction of solar flares from sdo images using deep learning. *Advances in Space Research* **67**(8), 2544–2557 (2021) <https://doi.org/10.1016/j.asr.2021.01.042>
- [29] Yi, K., Moon, Y.-J., Lim, D., Park, E., Lee, H.: Visual explanation of a deep learning solar flare forecast model and its relationship to physical parameters. *The Astrophysical Journal* **910**(1), 8 (2021)
- [30] Upendran, V., Cheung, M.C.M., Hanasoge, S., Krishnamurthi, G.: Solar Wind Prediction Using Deep Learning. *Space Weather* **18**(9), 02478 (2020) <https://doi.org/10.1029/2020SW002478> [arXiv:2006.05825](https://arxiv.org/abs/2006.05825) [astro-ph.SR]
- [31] Brown, E.J.E., Svoboda, F., Meredith, N.P., Lane, N., Horne, R.B.: Attention-Based Machine Vision Models and Techniques for Solar Wind Speed Forecasting Using Solar EUV Images. *Space Weather* **20**(3), 2021–002976 (2022) <https://doi.org/10.1029/2021SW00297610.1002/essoar.10508581.1>
- [32] Raju, H., Das, S.: Cnn-based deep learning model for solar wind forecasting. *Solar Physics* **296**(9), 134 (2021)
- [33] Bernoux, G., Brunet, A., Buchlin, É., Janvier, M., Sicard, A.: Forecasting the Geomagnetic Activity Several Days in Advance Using Neural Networks Driven by Solar EUV Imaging. *Journal of Geophysical Research (Space Physics)* **127**(10), 2022–030868 (2022) <https://doi.org/10.1029/2022JA030868>
- [34] Chifu, I., Gafeira, R.: 3d solar coronal loop reconstructions with machine learning. *The Astrophysical Journal Letters* **910**(1), 10 (2021)
- [35] Rahman, S., Shin, S., Jeong, H.-J., Siddique, A., Moon, Y.-J., Park, E., Kang, J., Bae, S.-H.: Fast reconstruction of 3d density distribution around the sun based on the mas by deep learning. *The Astrophysical Journal* **948**(1), 21 (2023)
- [36] Szenicer, A., Fouhey, D.F., Munoz-Jaramillo, A., Wright, P.J., Thomas, R., Galvez, R., Jin, M., Cheung, M.C.M.: A deep learning virtual instrument for monitoring extreme uv solar spectral irradiance. *Science Advances* **5**(10), 6548 (2019) <https://doi.org/10.1126/sciadv.aaw6548> <https://www.science.org/doi/pdf/10.1126/sciadv.aaw6548>

- [37] Woods, T.N., Eparvier, F., Hock, R., Jones, A., Woodraska, D., Judge, D., Didkovsky, L., Lean, J., Mariska, J., Warren, H., et al.: Extreme ultraviolet variability experiment (eve) on the solar dynamics observatory (sdo): Overview of science objectives, instrument design, data products, and model developments. *The solar dynamics observatory*, 115–143 (2012)
- [38] Dos Santos, L.F., Bose, S., Salvatelli, V., Neuberg, B., Cheung, M.C., Janvier, M., Jin, M., Gal, Y., Boerner, P., Baydin, A.G.: Multichannel autocalibration for the atmospheric imaging assembly using machine learning. *Astronomy & Astrophysics* **648**, 53 (2021)
- [39] Kim, T., Park, E., Lee, H., Moon, Y.-J., Bae, S.-H., Lim, D., Jang, S., Kim, L., Cho, I.-H., Choi, M., Cho, K.-S.: Solar farside magnetograms from deep learning analysis of STEREO/EUVI data. *Nature Astronomy* **3**, 397–400 (2019) <https://doi.org/10.1038/s41550-019-0711-5>
- [40] Jeong, H.-J., Moon, Y.-J., Park, E., Lee, H.: Solar coronal magnetic field extrapolation from synchronic data with ai-generated farside. *The Astrophysical Journal Letters* **903**(2), 25 (2020)
- [41] Jeong, H.-J., Moon, Y.-J., Park, E., Lee, H., Baek, J.-H.: Improved ai-generated solar farside magnetograms by stereo and sdo data sets and their release. *The Astrophysical Journal Supplement Series* **262**(2), 50 (2022)
- [42] Howard, R.A., Moses, J., Vourlidis, A., Newmark, J., Socker, D.G., Plunkett, S.P., Korendyke, C.M., Cook, J., Hurley, A., Davila, J., *et al.*: Sun earth connection coronal and heliospheric investigation (secchi). *Space Science Reviews* **136**, 67–115 (2008)
- [43] Son, J., Cha, J., Moon, Y.-J., Lee, H., Park, E., Shin, G., Jeong, H.-J.: Generation of he i 1083 nm images from sdo aia images by deep learning. *The Astrophysical Journal* **920**(2), 101 (2021)
- [44] Salvatelli, V., Bose, S., Neuberg, B., dos Santos, L.F.G., Cheung, M., Janvier, M., Gunes Baydin, A., Gal, Y., Jin, M.: Using U-Nets to Create High-Fidelity Virtual Observations of the Solar Corona. arXiv e-prints, 1911–04006 (2019) <https://doi.org/10.48550/arXiv.1911.04006> arXiv:1911.04006 [astro-ph.SR]
- [45] Salvatelli, V., Dos Santos, L.F., Bose, S., Neuberg, B., Cheung, M.C., Janvier, M., Jin, M., Gal, Y., Baydin, A.G.: Exploring the limits of synthetic creation of solar euv images via image-to-image translation. *The Astrophysical Journal* **937**(2), 100 (2022)
- [46] Lim, D., Moon, Y.-J., Park, E., Lee, J.-Y.: Selection of three (extreme) ultraviolet channels for solar satellite missions by deep learning. *The Astrophysical Journal Letters* **915**(2), 31 (2021)

- [47] Galvez, R., Fouhey, D.F., Jin, M., Szenicer, A., Muñoz-Jaramillo, A., Cheung, M.C., Wright, P.J., Bobra, M.G., Liu, Y., Mason, J., *et al.*: A machine-learning data set prepared from the nasa solar dynamics observatory mission. *The Astrophysical Journal Supplement Series* **242**(1), 7 (2019)
- [48] Park, E., Moon, Y.-J., Lee, J.-Y., Lee, H., Lim, D., Shin, G., Kim, T., *et al.*: Generation of solar uv and euV images from sdo/hmi magnetograms by deep learning. *The Astrophysical Journal Letters* **884**(1), 23 (2019)
- [49] Felipe, T., Ramos, A.A.: Improved detection of far-side solar active regions using deep learning. *Astronomy & Astrophysics* **632**, 82 (2019)
- [50] Baso, C.D., Ramos, A.A.: Enhancing sdo/hmi images using deep learning. *Astronomy & Astrophysics* **614**, 5 (2018)
- [51] Dou, F., Xu, L., Ren, Z., Zhao, D., Zhang, X.: Super-resolution of Solar Magnetograms Using Deep Learning. *Research in Astronomy and Astrophysics* **22**(8), 085018 (2022) <https://doi.org/10.1088/1674-4527/ac78ce>
- [52] Scherrer, P.H., Bogart, R.S., Bush, R., Hoeksema, J.-a., Kosovichev, A., Schou, J., Rosenberg, W., Springer, L., Tarbell, T., Title, A., *et al.*: The solar oscillations investigation—michelson doppler imager. *The soho mission*, 129–188 (1995)
- [53] Okoh, D.I., Seemala, G.K., Rabiou, A.B., Uwamahoro, J., Habarulema, J.B., Aggarwal, M.: A Hybrid Regression-Neural Network (HR-NN) Method for Forecasting the Solar Activity. *Space Weather* **16**(9), 1424–1436 (2018) <https://doi.org/10.1029/2018SW001907>
- [54] Benson, B., Pan, W.D., Prasad, A., Gary, G.A., Hu, Q.: Forecasting solar cycle 25 using deep neural networks. *Solar Physics* **295**(5), 65 (2020) <https://doi.org/10.1007/s11207-020-01634-y>
- [55] Li, Q., Wan, M., Zeng, S.-G., Zheng, S., Deng, L.-H.: Predicting the 25th solar cycle using deep learning methods based on sunspot area data. *Research in Astronomy and Astrophysics* **21**(7), 184 (2021) <https://doi.org/10.1088/1674-4527/21/7/184>
- [56] Prasad, A., Roy, S., Sarkar, A., Chandra Panja, S., Narayan Patra, S.: Prediction of solar cycle 25 using deep learning based long short-term memory forecasting technique. *Advances in Space Research* **69**(1), 798–813 (2022) <https://doi.org/10.1016/j.asr.2021.10.047>
- [57] Bizzarri, I., Barghini, D., Mancuso, S., Alessio, S., Rubinetti, S., Taricco, C.: Forecasting the solar cycle 25 using a multistep bayesian neural network. *Monthly Notices of the Royal Astronomical Society* **515**(4), 5062–5070 (2022)
- [58] Asensio Ramos, A., Cheung, M.C.M., Chifu, I., Gafeira, R.: Machine learning in

solar physics. *Living Reviews in Solar Physics* **20**(1), 4 (2023) <https://doi.org/10.1007/s41116-023-00038-x> arXiv:2306.15308 [astro-ph.SR]

- [59] The SunPy Community, Barnes, W.T., Bobra, M.G., Christe, S.D., Freij, N., Hayes, L.A., Ireland, J., Mumford, S., Perez-Suarez, D., Ryan, D.F., Shih, A.Y., Chanda, P., Glogowski, K., Hewett, R., Hughitt, V.K., Hill, A., Hiware, K., Inglis, A., Kirk, M.S.F., Konge, S., Mason, J.P., Maloney, S.A., Murray, S.A., Panda, A., Park, J., Pereira, T.M.D., Reardon, K., Savage, S., Sipócz, B.M., Stansby, D., Jain, Y., Taylor, G., Yadav, T., Rajul, Dang, T.K.: The sunpy project: Open source development and status of the version 1.0 core package. *The Astrophysical Journal* **890**, 68 (2020) <https://doi.org/10.3847/1538-4357/ab4f7a>
- [60] Svalgaard, Leif: Solar activity – past, present, future. *J. Space Weather Space Clim.* **3**, 24 (2013) <https://doi.org/10.1051/swsc/2013046>
- [61] MacTaggart, D., Prior, C., Raphaldini, B., Romano, P., Guglielmino, S.L.: Direct evidence that twisted flux tube emergence creates solar active regions. *Nature Communications* **12**(6621) (2021) <https://doi.org/10.1038/s41467-021-26981-7>
- [62] Wang, Y.-M., Ko, Y.-K.: Observations of slow solar wind from equatorial coronal holes. *The Astrophysical Journal* **880**(2), 146 (2019) <https://doi.org/10.3847/1538-4357/ab2add>
- [63] He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16000–16009 (2022)
- [64] Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning*, pp. 1597–1607 (2020). PMLR
- [65] Feichtenhofer, C., Li, Y., He, K., *et al.*: Masked autoencoders as spatiotemporal learners. *Advances in neural information processing systems* **35**, 35946–35958 (2022)
- [66] Hoffmann, S., Lessig, C.: Atmodist: Self-supervised representation learning for atmospheric dynamics. *Environmental Data Science* **2**, 6 (2023)
- [67] Dash, A., Ye, J., Wang, G., Jin, H.: High resolution solar image generation using generative adversarial networks. *Annals of Data Science*, 1–17 (2022)
- [68] Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3202–3211 (2022)
- [69] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby,

- N.: An image is worth 16x16 words: Transformers for image recognition at scale. CoRR **abs/2010.11929** (2020) [2010.11929](#)
- [70] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. CoRR **abs/2103.14030** (2021) [2103.14030](#)
- [71] Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., Wang, M.: Swin-unet: Unet-like pure transformer for medical image segmentation. In: European Conference on Computer Vision, pp. 205–218 (2022). Springer
- [72] Patro, B.N., Namboodiri, V.P., Agneeswaran, V.S.: Spectformer: Frequency and attention is what you need in a vision transformer. arXiv preprint arXiv:2304.06446 (2023)
- [73] Xu, M., Xiong, Y., Chen, H., Li, X., Xia, W., Tu, Z., Soatto, S.: Long short-term transformer for online action detection. *Advances in Neural Information Processing Systems* **34**, 1086–1099 (2021)
- [74] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International Conference on Machine Learning, pp. 2256–2265 (2015). PMLR
- [75] Price, I., Sanchez-Gonzalez, A., Alet, F., Ewalds, T., El-Kadi, A., Stott, J., Mohamed, S., Battaglia, P., Lam, R., Willson, M.: Gencast: Diffusion-based ensemble forecasting for medium-range weather. arXiv preprint arXiv:2312.15796 (2023)
- [76] Rühling Cachay, S., Zhao, B., Joren, H., Yu, R.: Dyffusion: A dynamics-informed diffusion model for spatiotemporal forecasting. *Advances in Neural Information Processing Systems* **36** (2024)
- [77] Ramunno, F.P., Hackstein, S., Kinakh, V., Drozdova, M., Quétant, G., Csillaghy, A., Voloshynovskiy, S.: Solar synthetic imaging: Introducing denoising diffusion probabilistic models on sdo/aia data. *Astronomy & Astrophysics* **686**, 285 (2024)
- [78] Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., et al.: Graphcast: Learning skillful medium-range global weather forecasting. arXiv preprint arXiv:2212.12794 (2022)
- [79] Lang, S., Chantry, M., team, T.A.: Enter the ensembles. <https://www.ecmwf.int/en/about/media-centre/aifs-blog/2024/enter-ensembles>. [Online; accessed September 26 2024] (2024)
- [80] Stanczuk, J.P., Batzolis, G., Deveney, T., Schönlieb, C.-B.: Diffusion models encode the intrinsic dimension of data manifolds. In: Forty-first International

- [81] Lippe, P., Veeling, B., Perdikaris, P., Turner, R., Brandstetter, J.: Pde-refiner: Achieving accurate long rollouts with neural pde solvers. *Advances in Neural Information Processing Systems* **36** (2024)
- [82] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695 (2022)