

CORONA-FIELDS: Leveraging Foundation Models for Classification of Solar Wind Phenomena

Daniela Martin^{1*}, Jinsu Hong^{2*}, Connor O'Brien³, Valmir P. Moraes Filho⁴, Jasmine R. Kobayashi⁵, Evangelia Samara⁶, Joseph Gallego⁷

¹University of Delaware

²Georgia State University

³Boston University

⁴Catholic University of America

⁵Southwest Research Institute

⁶NASA Goddard Space Flight Center

⁷Drexel University

dmartin@udel.edu, jhong36@gsu.edu, obrionco@bu.edu, moraesfilho@cua.edu, jasmine.kobayashi@swri.org, evangelia.samara@nasa.gov, jg3959@drexel.edu

Abstract

Space weather at Earth, driven by the solar activity, poses growing risks to satellites around our planet as well as to critical ground-based technological infrastructure. Major space weather contributors are the solar wind and coronal mass ejections whose variable density, speed, temperature, and magnetic field make the automated classification of those structures challenging. In this work, we adapt a foundation model for solar physics, originally trained on Solar Dynamics Observatory imagery, to create embeddings suitable for solar wind structure analysis. These embeddings are concatenated with the spacecraft position and solar magnetic connectivity encoded using Fourier features which generates a neural field-based model. The full deep learning architecture is fine-tuned bridging the gap between remote sensing and in situ observations. Labels are derived from Parker Solar Probe measurements, forming a downstream classification task that maps plasma properties to solar wind structures. Although overall classification performance is modest, likely due to coarse labeling, class imbalance, and limited transferability of the pretrained model, this study demonstrates the feasibility of leveraging foundation model embeddings for in situ solar wind tasks. As a first proof-of-concept, it lays the groundwork for future improvements toward more reliable space weather predictions. The code and configuration files used in this study are publicly available to support reproducibility.

Introduction

The solar wind is the continuous outflow of charged particles from the Sun into the interplanetary space (Parker 1958). It is composed of a variety of structures that differ in density, speed, temperature, magnetic field, and other properties depending on their source region on the Sun. Particularly, the fast component of the solar wind, otherwise known as high speed streams, that originates from coronal holes

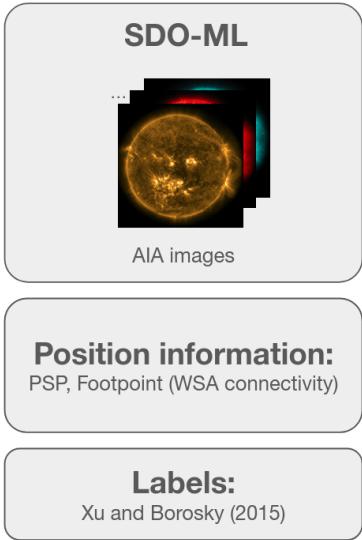
on the Sun, can cause medium- to large-scale geomagnetic storms (Richardson, Cliver, and Cane 2001; Echer, Tsurutani, and Gonzalez 2013). Such storms can disrupt satellites and pose radiations risks to astronauts. Within the solar wind, more extreme and transient phenomena frequently propagate such as coronal mass ejections (CMEs). CMEs are large expulsions of plasma and magnetic field from the Sun's corona (Forbes 2000; Webb and Howard 2012) that produce variations in the solar wind and, in turn, cause variations in the Earth's magnetic field as well (Richardson, Cliver, and Cane 2000). CMEs can cause large-scale geomagnetic storms which cannot only severely damage satellites (Baruah et al. 2024) and pose even higher risks than high speed streams to astronauts' life, but they can also disrupt terrestrial power grids (Hamrin et al. 2023) leading to significant economic consequences (Eastwood et al. 2017). Reliably forecasting both solar wind and CMEs remains highly challenging as our understanding of their structure and evolution is still incomplete (Owens et al. 2022).

Classifying solar wind structures (from now on when referring to solar wind we mean solar wind and CMEs together), is central to space weather research as solar wind structures originating from different sources (e.g., coronal holes, streamer belts, active regions, etc.) have distinct plasma and magnetic properties that determine their geo-effectiveness (Borovsky 2018). Being able to distinguish and predict these structures is key to reliably forecast geomagnetic activity and mitigate risks related to satellites, power grids, and astronauts. Several studies have attempted to predict solar wind properties directly from solar disk images. For example, Upendran et al. (2020) used EUV images to forecast wind speed, while Lin et al. (2023) employed magnetograms. More recently, Wang et al. (2025) predicted solar-wind-driven geomagnetic activity from EUV images. While these studies show promise, they typically rely on physics-based preprocessing of solar images before using them in neural networks. This step introduces uncertainty

*These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Model input



Model architecture

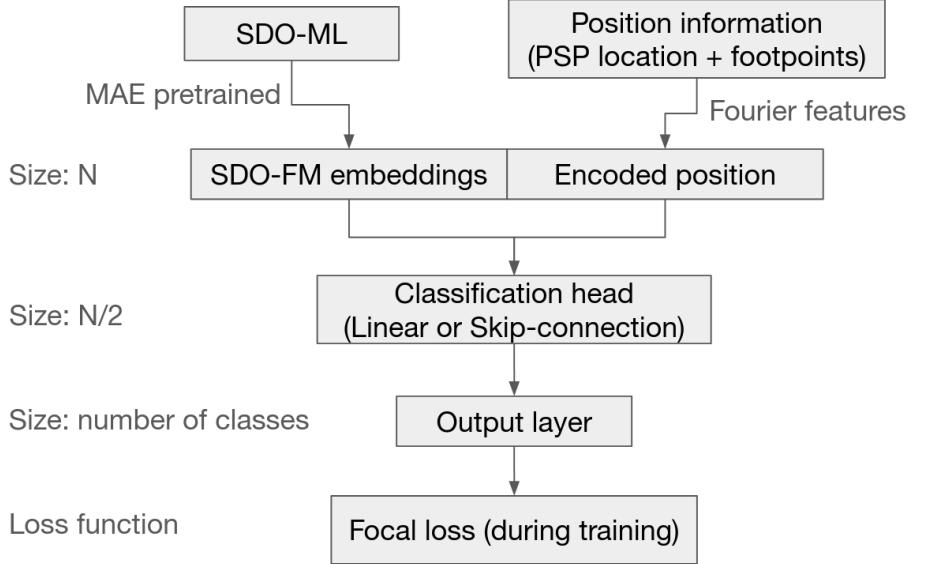


Figure 1: Architecture of the proposed model. The framework combines two input streams: (i) solar images processed by a pretrained MAE backbone (SDO-FM) to obtain an embedding layer, and (ii) positional information from the Parker Solar Probe and its magnetic footpoints, encoded through Fourier features. The embeddings and encoded positions are concatenated and passed through the classification head (either a simple linear head or a skip-connection variant), producing logits over the available solar wind classes. During training, predictions are compared with ground-truth labels using the focal loss function.

from model choices and may discard valuable information contained in the raw data.

Recent progress in self-supervised learning has enabled the training of foundation models, namely, large neural networks pretrained on raw, unlabeled data, learn general-purpose representations useful for many tasks (Dariset al. 2023). Notable examples include SimCLR (Chen et al. 2020), DINO (Caron et al. 2021), and MAE (He et al. 2022) foundation models. Specifically for MAE, it learns compact latent embeddings by reconstructing masked portions of the input. These embeddings can then be reused in downstream applications. MAE has already demonstrated strong transfer performance in fields like Earth observation (Allen et al. 2023; Gallego-Mejia et al. 2023), motivating their use in solar physics.

In this work, our goal is to design a deep learning framework to connect solar imagery with in situ solar wind observations (see overall architecture in Figure 1). We leverage the SDO Foundation Model (SDO-FM) (Walsh et al. 2024), a pretrained masked autoencoder (MAE) trained on solar images as a foundation model backbone for the downstream task of solar wind classification. To enhance representational power, we incorporate a neural field-based head architecture that encodes positional information in a high-frequency space, exploring both simple linear heads and more expressive variants based on ResNet (He et al. 2016) and NeRF (Mildenhall et al. 2021) skip connections. Recent work has highlighted the importance of enriching models with positional information. Incorporating this idea into our setup enables solar wind classification to take into account the geo-

metric context of the heliosphere, beyond what is available in the image embeddings alone. Solar wind structures are labeled using in situ plasma measurements mapped to the Xu and Borovsky (2015) scheme. Although overall classification performance is modest, the model appears to capture meaningful structure in the solar wind data, suggesting that foundation model embeddings can encode relevant physical information. Taken together, our results demonstrate that combining foundation models with in situ labeling schemes offers a first step toward effectively linking solar imagery with heliospheric plasma data, being a nontrivial step toward unifying solar physics and space weather forecasting.

Observational datasets

To close the gap in our knowledge of the solar wind, the Parker Solar Probe (PSP) (Fox et al. 2016; Szabo 2018) was launched in August 2018. One of the main objectives of this mission is the measurement of the solar wind properties at distances closer to the Sun than any other mission has achieved before (as close as ≈ 9 Solar Radii (R_{Sun}) or ≈ 0.04 Astronomical Units (AU)). That makes it the only spacecraft that has recorded in situ properties of the solar wind at an extended radial range within the inner heliosphere (≈ 0.04 AU to ≈ 1.0 AU; see (Raouafi et al. 2023) for a summary of findings from PSP’s initial mission). PSP carries four instrument suites—FIELDS, SWEAP, WISPR, and IS \odot IS—designed to measure different aspects of the solar wind. In this study, we use the FIELDS instrument (Bale et al. 2016), which measures the in situ magnetic field com-

ponents and the SWEAP instrument (Kasper et al. 2016), which measures the properties of the solar wind plasma.

In order to try to connect PSP in situ measurements to their solar source regions, we use remote-sensing images from the Solar Dynamics Observatory (SDO) (Pesnell, Thompson, and Chamberlin 2012). SDO was launched in 2010 and observes the solar atmosphere with a suite of instruments, including the Atmospheric Imaging Assembly (AIA) (Lemen et al. 2012) and the Helioseismic and Magnetic Imager (HMI) (Schou et al. 2012). AIA acquires 10-channel images of the solar disk (the full visible surface of the Sun as seen from Earth) at two ultraviolet (UV), seven extreme ultraviolet (EUV), and one visible wavelength (Lemen et al. 2011), while HMI observes the solar disk to derive global photospheric velocity, intensity, and magnetic field measurements (Hoeksema et al. 2014).

Although AIA and HMI provide high-resolution 4096×4096 images, such data is not yet well-suited for direct application in multi-modal learning with machine learning or deep learning algorithms (Jarolim et al. 2025). To address this, Galvez et al. (2019) introduced the SDOML dataset, a clean and machine-learning-oriented version of SDO observations downscaled to 512×512 , resulting in the reduction of computational demands. The SDOML dataset was designed to provide solar physicists and data scientists with a standardized, machine-learning-ready collection of SDO observations. To achieve this, raw Level-1 data from the three instruments on board SDO were first corrected to remove known instrumental artifacts such as flat-field irregularities, exposure variations, and bad pixels. Spatial and temporal downsampling was then applied to reduce the very high resolution of the original data to a more computationally manageable scale, while still preserving structures of interest for scientific and machine learning tasks. All image sequences were temporally aligned and co-registered across wavelengths to ensure that simultaneous multi-channel observations correspond to the same solar disk region and time frame, which is critical given the varying cadences of SDO instruments. Metadata, including helioprojective coordinates, solar ephemeris, and instrumental context, were retained in standardized formats, allowing downstream tasks to connect observations across physical domains. The resulting dataset therefore eliminates the need for extensive preprocessing at the user level, offering “analysis-ready” images where instrumental and observational inconsistencies have already been corrected. By leveraging SDOML, we can directly use compact embeddings from foundation models without repeating ad-hoc preprocessing pipelines (Walsh et al. 2024), which reduces uncertainty and ensures reproducibility across experiments. We use images from SDOML provided at a cadence of 12 minutes, matching the downsampled format described in (Galvez et al. 2019), sufficient to capture large-scale coronal structures relevant to solar wind origin while keeping the dataset computationally manageable.

Models

Backbone: SDO-FM

Several approaches have been proposed to build foundation models with SDO data. The SDO-FM model (Walsh et al. 2024) utilizes SDOML and leverages Masked Autoencoder (MAE) (He et al. 2022) and the Nouveau Variational Autoencoder (NVAE) (Vahdat and Kautz 2020) to construct the backbone. We build upon the SDO-FM (Walsh et al. 2024), a large-scale masked autoencoder (MAE) trained on the SDOML dataset from 2010–2023 using TPU clusters. This model serves as a pretrained backbone that encodes rich spatiotemporal representations of solar structures. During pretraining, random patches of the input images are masked, and the encoder learns to produce a compact latent representation that allows the decoder to reconstruct the missing regions. This self-supervised approach enables MAE to capture meaningful spatial and structural features of the solar corona (He et al. 2022; Walsh et al. 2024). The output embedding layer from the encoder is used with a custom head neural network to do the downstream task of solar wind classification.

The SDO-FM model used in this work was previously trained with AIA data from the SDOML dataset. Given the nature of its training, it is expected that it will primarily capture coronal intensity features rather than magnetic topological ones, which may limit its ability to identify potential magnetic drivers of the solar wind. The specific model was selected as it was already pretrained, allowing us to explore downstream solar wind classification while acknowledging its limitations.

Architecture of the Classification Head

To enhance the model’s representational power, we incorporate neural fields, positional information from the heliospheric PSP’s latitude and longitude and the photospheric footpoint of the magnetic field line that PSP is connected to based on Wang-Sheeley-Arge (WSA) (Arge and Pizzo 2000) predictions, through sinusoidal (Fourier) features into a high-dimensional spectral space (Rahimi and Recht 2007; Mildenhall et al. 2021). The mapping of a coordinate x is defined in Equation , where L is the number of frequency bands. This encoding captures both low- and high-frequency variations, providing richer information about the spacecraft’s position and its connection to the solar surface.

$$\gamma(x) = \left[\sin(2^0 \pi x), \cos(2^0 \pi x), \dots, \sin(2^{L-1} \pi x), \cos(2^{L-1} \pi x) \right],$$

The MAE embeddings and positional encodings are concatenated to form a unified representation, which serves as the input to the classification head. Two head variants were explored.

Linear head. The linear head consists of three fully connected layers with ReLU activations and dropout. The second and third layers use half the number of neurons of the preceding layer, respectively.

Skip-connection head. The skip-connection head follows the strategy of NeRF (Mildenhall et al. 2021), where intermediate layers receive both the hidden activations and the original input through skip connections every k layers. This allows the model to preserve fine-grained positional information across deeper transformations and enables richer interactions between pretrained embeddings and positional encodings. All the head’s layers have the same number of neurons.

Both architectures produce a final layer with four logits corresponding to the classes. Together, these two variants allow us to assess the trade-off between simplicity and expressive capacity in integrating positional information into the solar wind classification pipeline.

Training strategy

Our training follows a two-stage transfer learning and fine-tuning strategy:

1. **Transfer learning stage:** we initialize the backbone with the pretrained SDO-FM weights and freeze all its layers, training only the newly added classification head (and positional embedding layers) to adapt the representations to the solar wind classification task.
2. **Fine-tuning stage:** after convergence of the head, we unfreeze the entire network, including the MAE transformer backbone, allowing joint optimization of all parameters to refine the representations for in situ plasma classification.

Training is performed with the Adam optimizer (Kingma and Ba 2014) and a plateau learning rate scheduler. To handle class imbalance, we use the focal loss (Lin et al. 2017), defined as $FL(p_t) = -\alpha_t(1 - p_t)^\gamma \log(p_t)$, where p_t is the predicted probability for the true class, α_t balances class weights, and γ emphasizes hard-to-classify samples (Lin et al. 2017). This loss encourages the model to learn discriminative features even for underrepresented solar wind categories. As shown in Figure 1, the concatenated features are passed through the classification head, where the focal loss is applied during both training stages.

Experiment

In this section, we describe the preprocessing steps and evaluation protocols used to assess our model.

Preprocessing of PSP datasets

PSP measurements from the FIELDS and SWEAP instruments were first resampled to a common one-minute cadence via binning and averaging. SWEAP has a native cadence of approximately 25 seconds, while FIELDS measures at ≈ 3 ms; thus, each one-minute bin contains roughly 20,000 FIELDS and 2-3 SWEAP measurements. Minutes with missing data were linearly interpolated and flagged per instrument to allow later filtering if necessary. Additional parameters were calculated from PSP measurements to enrich the dataset. As mentioned earlier, the magnetic footprint of PSP on the solar surface and the corresponding solar wind travel time were obtained using the WSA model (Arge

and Pizzo 2000). These footpoints allow each measurement to be linked to the region of the Sun from which the measured solar wind parcel originated from. WSA was used for ease of calculation and as a baseline for image selection.

Solar wind structure labels were derived using a four-class segmentation scheme based on solar wind plasma conditions (Xu and Borovsky 2015). This classification partitions solar wind into fast plasma originating from coronal holes (Sheeley, Harvey, and Feldman 1976), slower wind plasma from the streamer belt region (Crooker et al. 2012), dense and cold sector reversal plasma forming the heliospheric current sheet (Susino et al. 2008), and ejecta which are linked to CMEs and magnetic clouds (Richardson, Cliver, and Cane 2000; Zhao, Zurbuchen, and Fisk 2009). Labels are based on proton temperature, proton specific entropy, and Alfvén speed, providing a notional link to the solar origin of each wind parcel. These four structural classes serve as the labels for the classification algorithm developed in this study. The structure of the algorithm is easily adaptable to other classification schemas; this one was chosen for its computational simplicity and notional link to the approximate area the solar wind originated from. It should be acknowledged, however, that this scheme presents challenges and has big uncertainties because of the fact that (1) the connectivity of the ejectas’ footpoints back to the Sun based on the WSA model are not reliable since the WSA model predicts only solar wind and no transient structures, and (2) the fact that the streamer belt and sector reversal categories often display overlapping properties resulting in some inconsistencies and overlapping in earlier classifications (Camporeale, Carè, and Borovsky 2017).

Connectivity of PSP measurements with SDOML images

In order to connect PSP measurements with an SDOML/AIA image, we assumed that WSA produces a good estimate of the solar wind at each PSP location. Then, for each such PSP location, we traced back to the photosphere the magnetic field line that the spacecraft was connected to. This is how each PSP measurement was paired with a SDOML/AIA image, producing a dataset of image-PSP position-footpoint-label. A single SDOML image can be associated with multiple PSP observations due to higher PSP dataset cadence when compared to the SDOML dataset. However, this does not introduce label noise into the dataset, since each pairing corresponds to a distinct heliospheric position and magnetic footpoint. In practice, this provides the model with complementary views of different source regions from the same solar image enhancing the diversity of spatial information available during training.

The training set includes data from April to December between 2019-2023. The validation set uses data from January to March between 2019-2022. The test set consists of data from January to March of 2023. Overall, the training, validation, and test sets comprise approximately 92%, 6%, and 2% of the dataset, respectively, ensuring temporal separation to avoid potential leakage due to solar rotation. Despite representing about 2% of the total dataset, the test split remains statistically meaningful given the overall dataset

size of nearly one million samples, each composed of 10-channel AIA imagery. This corresponds to over 13,000 test instances, which provides a sufficiently large sample to evaluate model generalization with low variance in performance metrics. The choice of using approximately three months of data for testing was deliberate: by reserving a contiguous temporal segment, the evaluation avoids data leakage across solar rotations and ensures that the test period is unseen both spatially and temporally. This design prioritizes independence and realism over uniform proportional splitting, which is critical in temporally correlated heliophysics datasets.

Table 1 summarizes the dataset partitions and class distributions. Figure 2 shows the inherently imbalanced nature of the data: streamer belt wind and sector reversal dominate, while coronal hole and ejecta are underrepresented, motivating the use of focal loss during training.

Training setup and evaluation

To evaluate model performance in the imbalanced multi-class classification task, we report accuracy, precision, and F1-score, all computed per class and averaged with a macro strategy. To mitigate class imbalance, we tuned the α and γ parameters of the focal loss. Models were trained for 50 epochs with a batch size of 32, and hyperparameters (including learning rate, weight decay, and scheduler) were selected based on focal loss validation. The full hyperparameter search space is summarized in Table 2. All experiments were implemented in PyTorch (FP32) and executed on Google Cloud Platform using four c2-standard-8 VM instances (8 vCPUs, 32 GB RAM) with 10 TB SSD storage and four NVIDIA A100 GPUs. The full codebase supporting this work is available at <https://github.com/spaceml-org/CORONA-FIELDS>.

Results

Figure 3 compares training loss across three configurations. The model with random initialization of the MAE backbone exhibits consistently higher loss across all epochs compared to the other two configurations. This outcome is expected, as leveraging a pretrained backbone provides prior knowledge about the image domain and avoids retraining it from scratch on potentially massive datasets, thereby facilitating faster and more effective optimization. Among the models with lower loss, the configuration with a fine-tuned pretrained backbone achieves the best performance, indicating that adapting pretrained features with new task-specific information leads to more effective learning. Overall, the comparison highlights the value of pretraining for stable and efficient optimization. Training was stopped early based on validation loss, as part of an early-stopping criterion to prevent overfitting once performance on the validation set plateaued.

In addition to evaluating training loss, we also analyzed how incorporating positional information affects the structure of the embedding space. After adding positional encoding, the embeddings are projected into a lower-dimensional space, as visualized in the t-SNE plots of embeddings with and without positional encoding (Figure 4). While the different types of solar wind structures are not fully separable

in either case, adding positional information reveals interesting patterns and a clearer organization of the embedding space, reflecting a dimensionality reduction induced by the positional encoding.

Figure 5 illustrates randomly selected correct predictions of the coronal hole, sector reversal, and streamer belt classes. For the coronal hole cases, the WSA-predicted magnetic footpoints lie in dark regions in the AIA images (especially visible in the 193 Å), which are coronal holes. For the streamer belt cases, the magnetic footpoints lie in the brighter, more dense regions of the Sun, characteristic of streamers. In the sector reversal case, the magnetic footpoints should lie around the area of the heliospheric current sheet.

Table 3 summarizes the performance of the evaluated architectures under three training regimes: training from scratch, transfer learning, and fine-tuning. The linear head achieves higher F1 when trained from scratch, while the skip-connection head performs best under fine-tuning. Overall, the linear head yields the highest accuracy across all settings. The model struggles to distinguish between streamer belt and sector reversal plasma classes, reflecting feature overlap and limitations of the Xu and Borovsky (2015) labeling scheme (Camporeale, Carè, and Borovsky 2017). Additionally, due to the scarcity of ejecta samples and WSA capabilities, the model does not learn a reliable representation for this class.

Although overall accuracy and F1 scores remain modest, the model appears to learn meaningful information from the AIA images. It seems to integrate features of the solar surface with the position of PSP (which is not in the field of view of the images) and the spacecraft's magnetic footpoints in the photosphere to inform its predictions. This suggests that even with limited performance, the pretrained embeddings capture relevant solar structures, highlighting the potential of foundation models for downstream in situ tasks.

Figure 6 shows an example of one SDOML/AIA 193 Å image and its reconstructions by the pretrained MAE and the fully fine-tuned model. Fine-tuning enhanced contrast, making active regions brighter and coronal holes darker compared to the pretrained MAE output, highlighting the model's adaptation to task-specific features.

Discussion and limitations

A primary factor limiting performance is the labeling scheme itself. The Xu and Borovsky (2015) scheme provides a simple empirical segmentation, but its use of fixed thresholds introduces ambiguity near transitional regimes where plasma properties overlap, a common limitation of threshold-based classification methods in the continuously varying solar wind environment. This ambiguity is particularly evident between the streamer belt and sector reversal plasma classes, whose boundaries are not sharply defined, and is further compounded by the scarcity of ejecta samples, preventing the model from learning a robust representation for that class. As a result, overall classification accuracy remains modest ($\approx 30\%$). The high similarity between the streamer belt and sector reversal classes suggests that some apparent misclassifications may not reflect a lack of

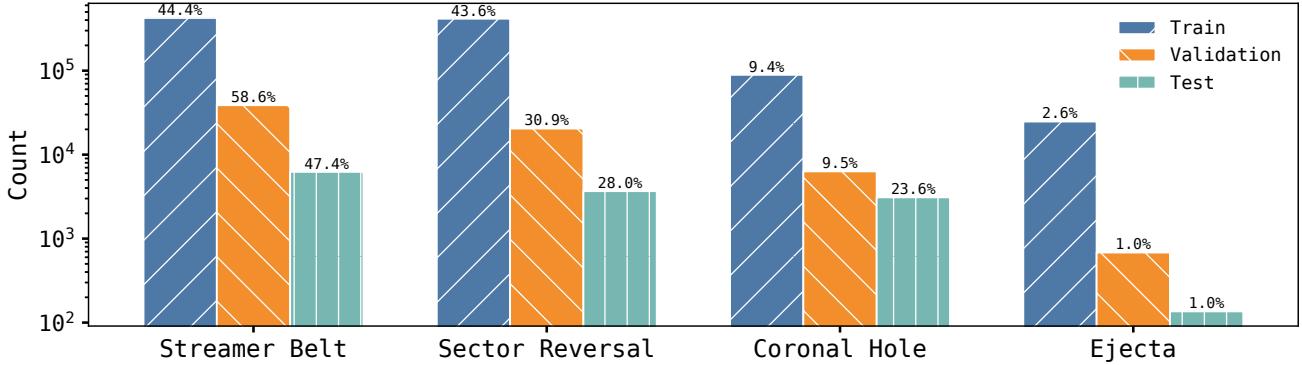


Figure 2: Distribution of each solar wind class across each dataset.

Set	Total	Streamer Belt	Sector Reversal	Coronal Hole	Ejecta
Train	953,821	415,870	423,960	89,206	24,785
Validation	66,245	38,799	20,444	6,319	683
Test	13,148	6,235	3,675	3,102	136

Table 1: Dataset partition and class distribution. Each sample corresponds to a 10-channel AIA image and its associated plasma parameters.

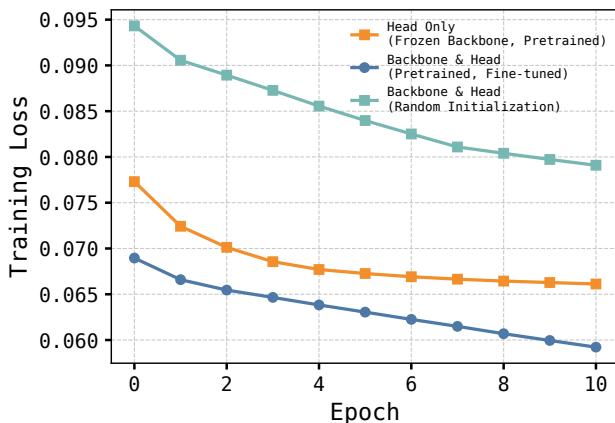


Figure 3: Training loss curves for a linear head with skip-connection. Three model configurations are compared: (1) head only, with a frozen, pretrained backbone, (2) backbone & head with a pretrained backbone fine-tuned during training, and (3) backbone & head with random initialization.

learned structure, but rather the intrinsic ambiguity of the labels. The model exhibits a tendency to predict the streamer belt category, which is expected given its high prevalence in the dataset and the intrinsic similarity with the sector reversal class. While combining the high prevalence classes has not been explicitly tested, it indicates that the model could be capturing meaningful relationships between coronal features and *in situ* plasma measurements beyond merely predicting the majority class.

Another factor potentially limiting performance is the backbone itself. The SDO-FM model used here was pretrained exclusively on AIA images, capturing primarily coronal intensity features rather than magnetic topology. As a result, the model may have limited ability to fully represent the magnetic drivers of the solar wind. While this choice was dictated by the availability of pretrained foundation models, it highlights an avenue for future work where incorporating HMI magnetograms could improve downstream classification.

The mapping of PSP measurements to photospheric footpoints via the WSA model is inherently approximate, and the subsequent resampling of data from 8-hour to 1-minute cadence may introduce interpolation artifacts and noise. Also, as mentioned earlier, the WSA backmapping does not apply for the ejecta category. These approximations affect the fidelity of the positional encoding, which could in turn limit the discriminative power of the head architectures.

Overall, while this model is not intended as an operational or immediately deployable system, it provides a valuable proof of concept and diagnostic tool. Its limited performance helps reveal where current assumptions, such as coarse threshold-based labeling, fall short, offering insights that can guide future developments in both modeling and data annotation. In this sense, the present work contributes

Parameter	Search Space
Head Type	Linear layers, Linear layers with skip-connection
Optimizer	Adam
Hidden Layer Size	64, 128, 256, 512, 1024
Learning Rate	$1 \times 10^{-5}, 1 \times 10^{-6}, 1 \times 10^{-7}, 1 \times 10^{-8}$
Weight Decay	$3 \times 10^{-4}, 1 \times 10^{-4}, 1 \times 10^{-3}$
Scheduler	Cosine Annealing, Reduce on Plateau
Loss Function	Cross-Entropy, Focal Loss
Focal Loss α	[0.45, 0.30, 0.15, 0.10], [0.45, 0.35, 0.10, 0.10]
Focal Loss γ	2, 3
Sampling Strategy	No modification, Under-sampling

Table 2: Hyperparameter search space for the solar wind classification task.

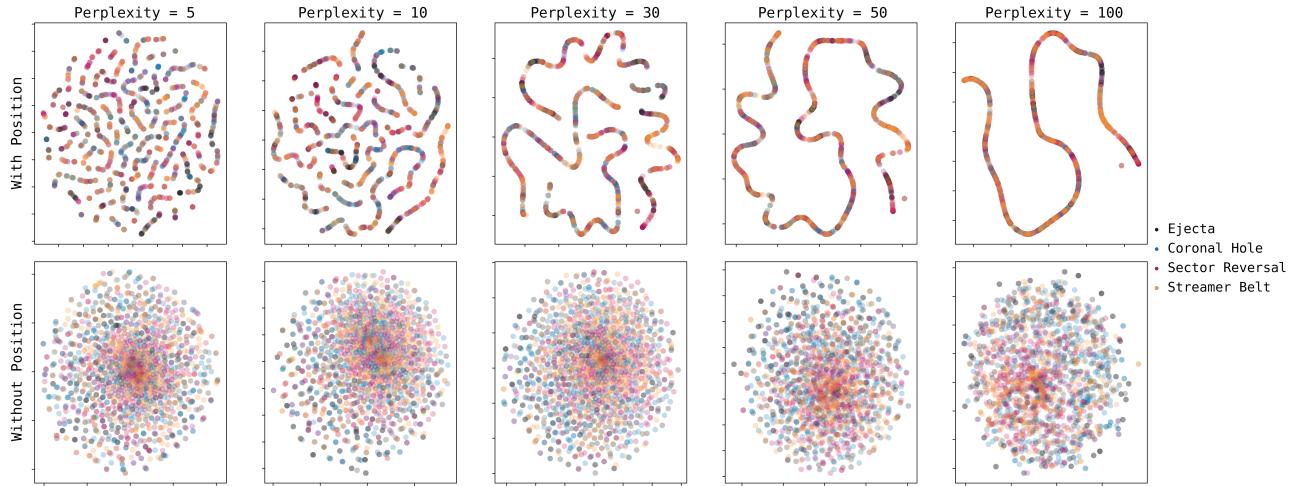


Figure 4: t-SNE on a randomly selected balanced set of embeddings produced by SDO-FM with PSP’s position included (top row) and without PSP’s position included (bottom row) at five different perplexities. With positional information included, the embeddings are projected into a lower-dimensional space, as evidenced by the strip-like structure in the t-SNE visualization.

not only a methodology, but also a critical reflection on existing empirical schemes in heliophysics.

Conclusion

In this work, we adapted a foundation model for solar physics, combining MAE embeddings with neural field-based head architectures, to bridge remote-sensing imagery from the SDO with in situ plasma measurements from PSP for the classification of solar wind structures. Although the classification performance remains modest, the results demonstrate a proof of concept that positional encodings of spacecraft location and magnetic connectivity can enrich pretrained embeddings for downstream heliospheric tasks.

This study demonstrates that foundation models can serve as a bridge between solar imagery and heliospheric in situ observations, an essential and nontrivial step toward unified models for solar physics and space weather forecasting. The presented pipeline represents a first step in exploring how foundation model embeddings can be extended from solar remote sensing to in situ applications. While current results are limited, they reveal that embeddings learned from coronal imagery capture meaningful physical informa-

tion that can inform solar wind classification. This connection between modalities highlights a promising pathway toward more robust, interpretable, and data-driven approaches to space weather analysis.

Future Work

Several avenues could be explored to address the limitations identified in this study. First, improving the labeling scheme by incorporating hybrid or probabilistic methods could reduce ambiguity near transitional regimes and better capture the diversity of solar wind structures. Second, extending the foundation model pretraining to include additional magnetic field data, such as HMI magnetograms, could enhance the transferability of embeddings for in situ tasks. Third, larger and more balanced test and validation splits, along with metrics tailored for class imbalance, would provide a more robust assessment of model performance. Finally, optimizing data storage and loading pipelines could improve training scalability for larger datasets and more complex models, enabling more comprehensive experimentation.

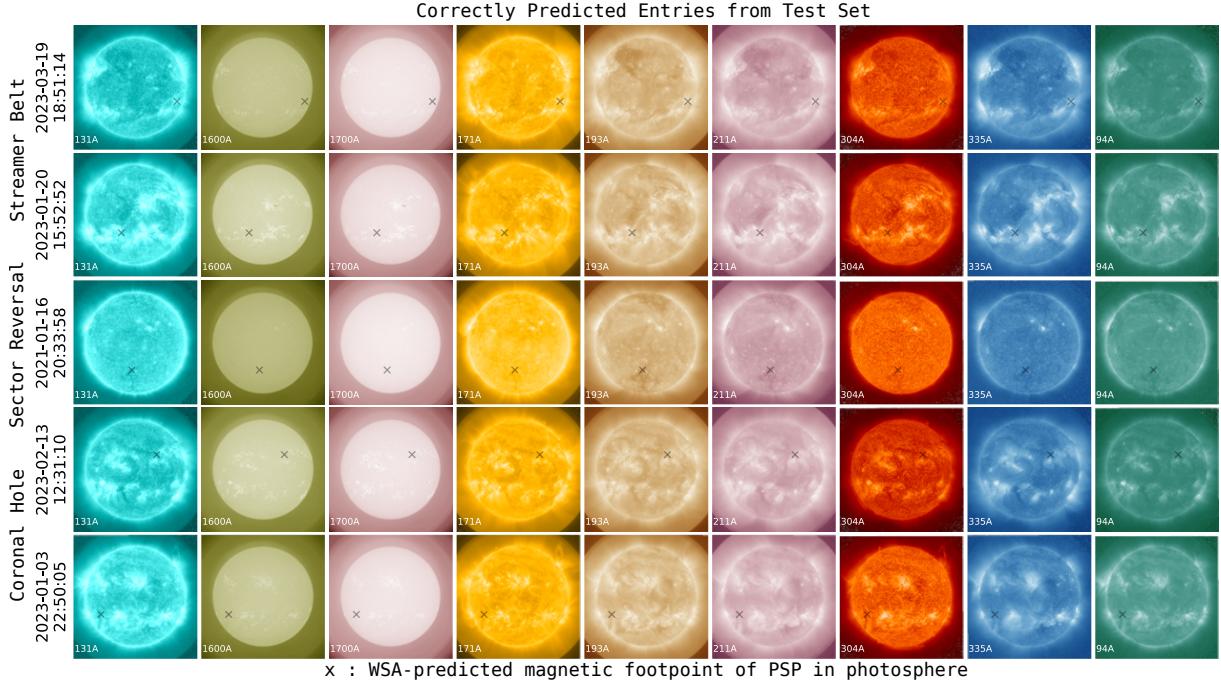


Figure 5: Examples of correct predictions on the test set for three solar wind classes: coronal hole, sector reversal, and streamer belt. Each column shows an SDOML/AIA image, with crosses indicating the predicted magnetic footprints from PSP.

Head Type	Training Strategy	Accuracy	Precision	Recall	F1-score
Skip-connection	Head Only (Frozen Backbone, Pretrained)	0.305	0.261	0.305	0.242
	Backbone & Head (Pretrained, Fine-tuned)	0.345	0.304	0.345	0.308
	Backbone & Head (Random Initialization)	0.291	0.296	0.291	0.288
Linear	Head Only (Frozen Backbone, Pretrained)	0.314	0.287	0.314	0.288
	Backbone & Head (Pretrained, Fine-tuned)	0.325	0.293	0.325	0.298
	Backbone & Head (Random Initialization)	0.320	0.327	0.320	0.314

Table 3: Test set performance of linear and skip-connection heads under three training strategies: frozen pretrained backbone, fine-tuned pretrained backbone, and random initialization. Metrics include Accuracy, Precision, Recall, and F1-score.

Acknowledgements

This work is a research product of Heliolab (heliolab.ai), an initiative of the Frontier Development Lab (FDL.ai). FDL is a public-private partnership between NASA, Trillium Technologies (trillium.tech), and commercial AI partners including Google Cloud and NVIDIA. Heliolab was designed, delivered, and managed by Trillium Technologies Inc., a research and development company focused on intelligent systems and collaborative communities for Heliophysics, planetary stewardship and space exploration. We gratefully acknowledge Google Cloud for extensive computational resources, and NVIDIA Corporation for access to DGX Cloud, enabled through NVIDIA and VMware. This material is based upon work supported by NASA under award No. 80GSFC23CA040. Any opinions, findings, and conclusions or recommendations expressed are those of the author(s) and do not necessarily reflect the views of the National Aeronautics and Space Administration.

References

Allen, M.; Dorr, F.; Gallego-Mejia, J. A.; Martínez-Ferrer, L.; Jungbluth, A.; Kalaitzis, F.; and Ramos-Pollán, R. 2023. Fewshot learning on global multimodal embeddings for earth observation tasks. *arXiv preprint arXiv:2310.00119*.

Arge, C. N.; and Pizzo, V. J. 2000. Improvement in the prediction of solar wind conditions using near-real time solar magnetic field updates. *Journal of Geophysical Research: Space Physics*, 105(A5): 10465–10479.

Bale, S. D.; Goetz, K.; Harvey, P. R.; Turin, P.; Bonnell, J. W.; Dudok de Wit, T.; Ergun, R. E.; MacDowall, R. J.; Pulupa, M.; Andre, M.; Bolton, M.; Bougeret, J.-L.; Bowen, T. A.; Burgess, D.; Cattell, C. A.; Chandran, B. D. G.; Chaston, C. C.; Chen, C. H. K.; Choi, M. K.; Connerney, J. E.; Cranmer, S.; Diaz-Aguado, M.; Donakowski, W.; Drake, J. F.; Farrell, W. M.; Fergeau, P.; Fermin, J.; Fischer, J.; Fox, N.; Glaser, D.; Goldstein, M.; Gordon, D.; Hanson, E.; Harris, S. E.; Hayes, L. M.; Hinze, J. J.; Hollweg, J. V.; Horbury, T. S.; Howard, R. A.; Hoxie, V.; Jannet, G.; Karlsson, M.; Kasper, J. C.; Kellogg, P. J.; Kien, M.; Klimchuk, J. A.; Krasnoselskikh, V. V.; Krucker, S.; Lynch, J. J.; Maksimovic, M.; Malaspina, D. M.; Marker, S.; Martin, P.; Martinez-Oliveros, J.;

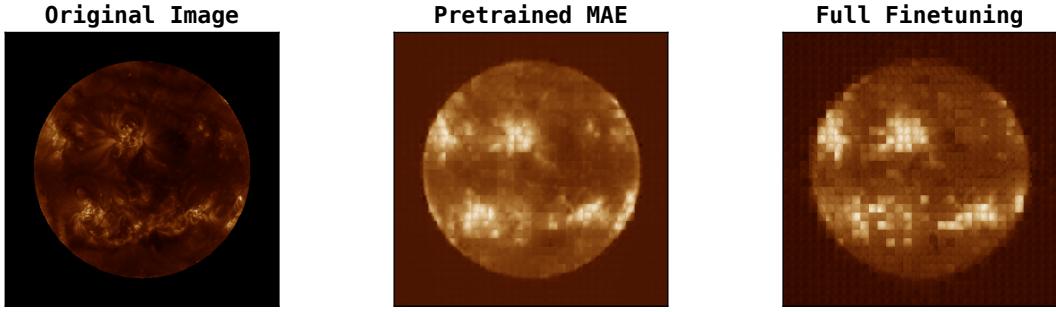


Figure 6: Example of image reconstruction using the MAE decoder. **Left:** original SDOML/AIA 193 Å image (2023-01-03 05:38:00). **Middle:** reconstruction with the pretrained MAE without fine-tuning. **Right:** reconstruction after fine-tuning with a linear classification head.

McCauley, J.; McComas, D. J.; McDonald, T.; Meyer-Vernet, N.; Moncuquet, M.; Monson, S. J.; Mozer, F. S.; Murphy, S. D.; Odom, J.; Oliverson, R.; Olson, J.; Parker, E. N.; Pankow, D.; Phan, T.; Quataert, E.; Quinn, T.; Ruplin, S. W.; Salem, C.; Seitz, D.; Sheppard, D. A.; Siy, A.; Stevens, K.; Summers, D.; Szabo, A.; Timofeeva, M.; Vaivads, A.; Velli, M.; Yehle, A.; Werthimer, D.; and Wygant, J. R. 2016. The FIELDS Instrument Suite for Solar Probe Plus: Measuring the Coronal Plasma and Magnetic Field, Plasma Waves and Turbulence, and Radio Signatures of Solar Transients. *Space Science Reviews*, 204(1-4): 49–82.

Baruah, Y.; Roy, S.; Sinha, S.; Palmerio, E.; Pal, S.; Oliveira, D. M.; and Nandy, D. 2024. The Loss of Starlink Satellites in February 2022: How Moderate Geomagnetic Storms Can Adversely Affect Assets in Low-Earth Orbit. *Space Weather*, 22(4): e2023SW003716.

Borovsky, J. E. 2018. The spatial structure of the oncoming solar wind at Earth and the shortcomings of a solar-wind monitor at L1. *Journal of Atmospheric and Solar-Terrestrial Physics*, 177: 2–11.

Camporeale, E.; Carè, A.; and Borovsky, J. E. 2017. Classification of solar wind with machine learning. *Journal of Geophysical Research: Space Physics*, 122(11): 10–910.

Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9650–9660.

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.

Crooker, N. U.; Antiochos, S. K.; Zhao, X.; and Neugebauer, M. 2012. Global network of slow solar wind. *Journal of Geophysical Research: Space Physics*, 117(A4): 2011JA017236.

Darcet, T.; Oquab, M.; Mairal, J.; and Bojanowski, P. 2023. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*.

Eastwood, J. P.; Biffis, E.; Hapgood, M. A.; Green, L.; Bisi, M. M.; Bentley, R. D.; Wicks, R.; McKinnell, L.; Gibbs, M.; and Burnett, C. 2017. The Economic Impact of Space Weather: Where Do We Stand? *Risk Analysis*, 37(2): 206–218.

Echer, E.; Tsurutani, B.; and Gonzalez, W. 2013. Interplanetary origins of moderate (- 100 nT; Dst- 50 nT) geomagnetic storms during solar cycle 23 (1996–2008). *Journal of Geophysical Research: Space Physics*, 118(1): 385–392.

Forbes, T. 2000. A review on the genesis of coronal mass ejections. *Journal of Geophysical Research: Space Physics*, 105(A10): 23153–23165.

Fox, N. J.; Velli, M. C.; Bale, S. D.; Decker, R.; Driesman, A.; Howard, R. A.; Kasper, J. C.; Kinnison, J.; Kusterer, M.; Lario, D.; Lockwood, M. K.; McComas, D. J.; Raouafi, N. E.; and Szabo, A. 2016. The Solar Probe Plus Mission: Humanity’s First Visit to Our Star. *Space Science Reviews*, 204(1-4): 7–48.

Gallego-Mejía, J. A.; Jungbluth, A.; Martínez-Ferrer, L.; Allen, M.; Dorr, F.; Kalaitzis, F.; and Ramos-Pollán, R. 2023. Exploring DINO: Emergent Properties and Limitations for Synthetic Aperture Radar Imagery. *arXiv preprint arXiv:2310.03513*.

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. 2019. A machine-learning data set prepared from the NASA solar dynamics observatory mission. *The Astrophysical Journal Supplement Series*, 242(1): 7.

Hamrin, M.; Schillings, A.; Opgenoorth, H.; Nesbit-Östman, S.; Krämer, E.; Araújo, J.; Baddeley, L.; Gunell, H.; Pitkänen, T.; Gjerloev, J.; and Barnes, R. J. 2023. Space Weather Disturbances in Non-Stormy Times: Occurrence of dB/dt Spikes During Three Solar Cycles. *Journal of Geophysical Research: Space Physics*, 128(10): e2023JA031804.

He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16000–16009.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Hoeksema, J. T.; Liu, Y.; Hayashi, K.; Sun, X.; Schou, J.; Couvidat, S.; Norton, A.; Bobra, M.; Centeno, R.; Leka, K. D.; Barnes, G.; and Turmon, M. 2014. The Helioseismic and Magnetic Imager (HMI) Vector Magnetic Field Pipeline: Overview and Performance. *Solar Physics*, 289(9): 3483–3530.

Jarolim, R.; Veronig, A.; Pötzl, W.; and Podladchikova, T. 2025. A deep learning framework for instrument-to-instrument translation of solar observation data. *Nature Communications*, 16(1): 3157.

Kasper, J. C.; Abiad, R.; Austin, G.; Balat-Pichelin, M.; Bale, S. D.; Belcher, J. W.; Berg, P.; Bergner, H.; Berthomier, M.; Bookbinder, J.; Brodu, E.; Caldwell, D.; Case, A. W.; Chandran, B. D. G.; Cheimets, P.; Cirtain, J. W.; Cranmer, S. R.; Curtis, D. W.; Daigneau, P.; Dalton, G.; Dasgupta, B.; DeTomaso, D.; Diaz-Aguado, M.; Djordjevic, B.; Donaskowski, B.; Effinger, M.; Florinski, V.; Fox, N.; Freeman, M.; Gallagher, D.; Gary, S. P.; Gauron, T.; Gates, R.; Goldstein, M.; Golub, L.; Gordon, D. A.; Gurnee, R.; Guth, G.; Halekas, J.; Hatch, K.; Heerikuisen, J.; Ho, G.; Hu,

Q.; Johnson, G.; Jordan, S. P.; Korreck, K. E.; Larson, D.; Lazarus, A. J.; Li, G.; Livi, R.; Ludlam, M.; Maksimovic, M.; McFadden, J. P.; Marchant, W.; Maruca, B. A.; McComas, D. J.; Messina, L.; Mercer, T.; Park, S.; Peddie, A. M.; Pogorelov, N.; Reinhart, M. J.; Richardson, J. D.; Robinson, M.; Rosen, I.; Skoug, R. M.; Slagle, A.; Steinberg, J. T.; Stevens, M. L.; Szabo, A.; Taylor, E. R.; Tiu, C.; Turin, P.; Velli, M.; Webb, G.; Whittlesey, P.; Wright, K.; Wu, S. T.; and Zank, G. 2016. Solar Wind Electrons Alphas and Protons (SWEAP) Investigation: Design of the Solar Wind and Coronal Plasma Instrument Suite for Solar Probe Plus. *Space Science Reviews*, 204(1-4): 131–186.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

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.; Gummin, M. A.; Smith, P.; Auker, G.; Jerram, P.; Pool, P.; Soufli, R.; Windt, D. L.; Beardsley, S.; Clapp, M.; Lang, J.; and Waltham, N. 2012. The Atmospheric Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO). *Solar Physics*, 275(1-2): 17–40.

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.; Gummin, M. A.; Smith, P.; Auker, G.; Jerram, P.; Pool, P.; Soufli, R.; Windt, D. L.; Beardsley, S.; Clapp, M.; Lang, J.; and Waltham, N. 2011. The Atmospheric Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO). *Solar Physics*, 275(1-2): 17–40.

Lin, R.; Luo, Z.; He, J.; Xie, L.; Hou, C.; and Chen, S. 2023. Prediction of solar wind speed by applying convolutional neural network to potential field source surface (PFSS) magnetograms. *ArXiv:2304.01234 [astro-ph, physics:physics]*.

Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.

Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.

Owens, M. J.; Chakraborty, N.; Turner, H.; Lang, M.; Riley, P.; Lockwood, M.; Barnard, L. A.; and Chi, Y. 2022. Rate of Change of Large-Scale Solar-Wind Structure. *Solar Physics*, 297(7): 83.

Parker, E. N. 1958. Dynamics of the Interplanetary Gas and Magnetic Fields. *The Astrophysical Journal*, 128: 664.

Pesnell, W. D.; Thompson, B. J.; and Chamberlin, P. C. 2012. *The Solar Dynamics Observatory (SDO)*, 3–15. New York, NY: Springer New York. ISBN 978-1-4614-3673-7.

Rahimi, A.; and Recht, B. 2007. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20.

Raouafi, N. E.; Matteini, L.; Squire, J.; Badman, S. T.; Velli, M.; Klein, K. G.; Chen, C. H. K.; Matthaeus, W. H.; Szabo, A.; Linton, M.; Allen, R. C.; Szalay, J. R.; Bruno, R.; Decker, R. B.; Akhavan-Tafti, M.; Agapitov, O. V.; Bale, S. D.; Bandyopadhyay, R.; Battams, K.; Berčič, L.; Bourouaine, S.; Bowen, T. A.; Cattell, C.; Chandran, B. D. G.; Chhiber, R.; Cohen, C. M. S.; D'Amicis, R.; Giacalone, J.; Hess, P.; Howard, R. A.; Horbury, T. S.; Jagalamudi, V. K.; Joyce, C. J.; Kasper, J. C.; Kinnison, J.; Laker, R.; Liewer, P.; Malaspina, D. M.; Mann, I.; McComas, D. J.; Niembro-Hernandez, T.; Nieves-Chinchilla, T.; Panasenco, O.; Pokorný, P.; Pusack, A.; Pulupa, M.; Perez, J. C.; Riley, P.; Rouillard, A. P.; Shi, C.; Stenborg, G.; Tenerani, A.; Verniero, J. L.; Viall, N.; Vourlidas, A.; Wood, B. E.; Woodham, L. D.; and Woolley, T. 2023. Parker Solar Probe: Four Years of Discoveries at Solar Cycle Minimum. *Space Science Reviews*, 219(1). Publisher: Springer Science and Business Media LLC.

Richardson, I.; Cliver, E.; and Cane, H. 2001. Sources of geomagnetic storms for solar minimum and maximum conditions during 1972–2000. *Geophysical Research Letters*, 28(13): 2569–2572.

Richardson, I. G.; Cliver, E. W.; and Cane, H. V. 2000. Sources of geomagnetic activity over the solar cycle: Relative importance of coronal mass ejections, high-speed streams, and slow solar wind. *Journal of Geophysical Research: Space Physics*, 105(A8): 18203–18213.

Schou, J.; Scherrer, P. H.; Bush, R. I.; Wachter, R.; Couvidat, S.; Rabbello-Soares, M. C.; Bogart, R. S.; Hoeksema, J. T.; Liu, Y.; Duvall, T. L.; Akin, D. J.; Allard, B. A.; Miles, J. W.; Raiden, R.; Shine, R. A.; Tarbell, T. D.; Title, A. M.; Wolfson, C. J.; Elmore, D. F.; Norton, A. A.; and Tomczyk, S. 2012. 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.

Sheeley, N. R.; Harvey, J. W.; and Feldman, W. C. 1976. Coronal holes, solar wind streams, and recurrent geomagnetic disturbances: 1973–1976. *Solar Physics*, 49(2): 271–278.

Susino, R.; Ventura, R.; Spadaro, D.; Vourlidas, A.; and Landi, E. 2008. Physical parameters along the boundaries of a mid-latitude streamer and in its adjacent regions. *Astronomy & Astrophysics*, 488(1): 303–310.

Szabo, A. 2018. Flying into the Sun. *Nature Astronomy*, 2(10): 829–829.

Upendran, V.; Cheung, M. C. M.; Hanasoge, S.; and Krishnamurthi, G. 2020. Solar Wind Prediction Using Deep Learning. *Space Weather*, 18(9): e2020SW002478.

Vahdat, A.; and Kautz, J. 2020. NVAE: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33: 19667–19679.

Walsh, J.; Gass, D. G.; Pollan, R. R.; Wright, P. J.; Galvez, R.; Kasmanoff, N.; Naradowsky, J.; Spalding, A.; Parr, J.; and Baydin, A. G. 2024. A Foundation Model for the Solar Dynamics Observatory. *arXiv preprint arXiv:2410.02530*.

Wang, T.; Luo, B.; Wang, J.; Ao, X.; Shi, L.; Zhong, Q.; and Liu, S. 2025. Forecasting of the Geomagnetic Activity for the Next 3 Days Utilizing Neural Networks Based on Parameters Related to Large-Scale Structures of the Solar Corona. *Space Weather*, 23(2): e2024SW004090.

Webb, D. F.; and Howard, T. A. 2012. Coronal mass ejections: Observations. *Living Reviews in Solar Physics*, 9(1): 1–83.

Xu, F.; and Borovsky, J. E. 2015. A new four-plasma categorization scheme for the solar wind. *Journal of Geophysical Research: Space Physics*, 120(1): 70–100. Publisher: American Geophysical Union (AGU).

Zhao, L.; Zurbuchen, T. H.; and Fisk, L. A. 2009. Global distribution of the solar wind during solar cycle 23: ACE observations. *Geophysical Research Letters*, 36(14): 2009GL039181.