
A Foundation Model for the Solar Dynamics Observatory

James Walsh
University of Cambridge

Daniel G. Gass
University of Central Lancashire

Raul Ramos Pollan
Universidad Industrial de Santander

Paul J. Wright
Dublin Institute for Advanced Studies

Richard Galvez
Pure Storage

Noah Kasmanoff
AE Studio

Jason Naradowsky
University of Tokyo

Anne Spalding
Trillium Technologies Inc

James Parr
Trillium Technologies Inc

Atılım Güneş Baydin
University of Oxford

Abstract

SDO-FM is a foundation model using data from NASA’s Solar Dynamics Observatory (SDO) spacecraft; integrating three separate instruments to encapsulate the Sun’s complex physical interactions into a multi-modal embedding space. This model can be used to streamline scientific investigations involving SDO by making the enormous datasets more computationally accessible for heliophysics research and enable investigations that require instrument fusion. We discuss four key components: an ingestion pipeline to create machine learning ready datasets, the model architecture and training approach, resultant embeddings & fine-tunable models, and finally downstream fine-tuned applications. A key component of this effort has been to include subject matter specialists at each stage of development; reviewing the scientific value and providing guidance for model architecture, dataset, and training paradigm decisions. This paper marks release of our pretrained models and embedding datasets, available to the community on Hugging Face and *sdo_fm.org*.

1 Introduction

NASA’s Solar Dynamics Observatory (SDO) [1] is a space based heliophysics laboratory, capturing details of the Sun’s atmosphere and photosphere across a wide variety of spatial and temporal scales, wavelengths and levels of activity across the solar cycle. It was launched in 2010 as part of NASA’s Living With a Star Program and has been instrumental in advancing the fields of observational heliophysics. It has captured the entirety of Solar Cycle 24 (2010 - 2020), and is currently observing the peak of Solar Cycle 25 (2021 - ongoing). The SDO consists of a suite of three instruments; the Helioseismic and Magnetic Imager (HMI) [1] which captures information on the Sun’s magnetic field at the photosphere, the Atmospheric Imaging Assembly (AIA) [2] which captures ultraviolet and soft X-ray radiation emitted from the Sun’s atmosphere in ten channels from 10-100 nm, and the Extreme-ultraviolet Variability Experiment (EVE) [3], which produces spectra for extreme ultraviolet irradiance reaching Earth’s atmosphere. This suite of instruments represents a large dataset suitable for use in training a foundation model, as well as a dataset with direct relevance to ongoing scientific research. The unique challenges and properties of the SDO dataset are outlined in Section 1.1.

The overall process for building foundation model for SDO, known as SDO-FM, is composed of four stages; (1) data preparation, (2) large foundation model (FM) training, (3) embedding extraction, and (4) fine-tuning or direct embedding usage for scientific validation cases. Collectively we denote the data preparation as effort completed under SDOML [4], a machine-learning dataset of SDO. Our models are based upon autoencoders, with training conducted under the objective of image reconstruction over the period beginning from satellite launch in 2010 to 2023. Once these models are trained, a compressed representation dataset is created from the embeddings by a full-pass over the encoder. The compressed representations are called direct embeddings and provide a helpful result as a set of available SDO features at around two-thousandths (0.002) the original size. Lastly, the direct embeddings as well as standard model fine-tuning, are used to conduct scientific validation through a validation harness which is used to check our results against past ML-based heliophysics approaches and to compare their computational expense.

1.1 Input Data

To properly represent the native SDO dataset in a format which is compatible with conventional machine learning architectures, the SDOML dataset [4] was used. This dataset is available at <http://sdoml.org>, and is comprised of 14 years of SDO data from all three instruments (HMI, AIA, and EVE). The high resolution and time cadence of AIA (4096 x 4096 pixels, 12 second cadence) and HMI (4096 x 4096) results in a database greater than 22 Pb in size on disk, comprised of multiple domains of image, magnetic flux, and spectra, as well as multiple image wavelengths.

To reduce the 14 years of data, SDOML v2.0 is offered at a fixed resolution and cadence (512x512 pixels every 12-minutes), with standardization and corrections. This includes transformations from Level-1 to 1.5 [5]. The size of the reduced dataset on disk of the 14-year period is roughly 10Tb, a more manageable state compared to the 22Pb source.

1.2 Scientific Direction of SDO-FM

The objective of this foundation model for heliophysics is to encode a representation of the phenomena observed by the SDO spacecraft. This includes key aspects of solar physics, such as the relationship between the solar magnetic field, eruptions (flaring), and extreme ultraviolet irradiance, as well as dynamics such as differential rotation, convection & magnetic processes, and observational factors such as instrumental degradation. Once physical phenomena are encoded in the embedding space, researchers can then develop fine-tuned downstream applications to explore how these phenomena impact the wider solar system - known as the Heliosphere. To achieve this ambitious goal, we begin with encoding known phenomena through a latent representation discussed further in section 3.1. The model is then validated by adapting the embeddings and comparing outcomes against published known results using classical machine learning methods. The classical comparisons were selected for applicability to the compute constraint problem: predicting the Earth's thermospheric density (F10.7), reconstructing missing channels, autocalibration of the AIA instrument, and the virtualization of the broken MEGS-A instrument, detailed in section 3.2.

2 Related Work

The idea for this approach came from previous efforts to encode the solar disk [6] and employ those embeddings to open heliophysics challenges, such as thermospheric density prediction [7]. It was evident from these investigations that there was merit to including SDO EUV data to improve those domain objectives, however these initial efforts were impeded by the high computational cost of training with large-scale image datasets, prompting the need for a foundational community asset that could be adapted to multiple heliophysics and space weather use-cases.

The precedent for employment of the MAE came from the proposal of MAEs as spatiotemporal learners [8] and the NASA IMPACT-IBM progress on a generalizable geospatial foundation model on remote sensing imagery [9].

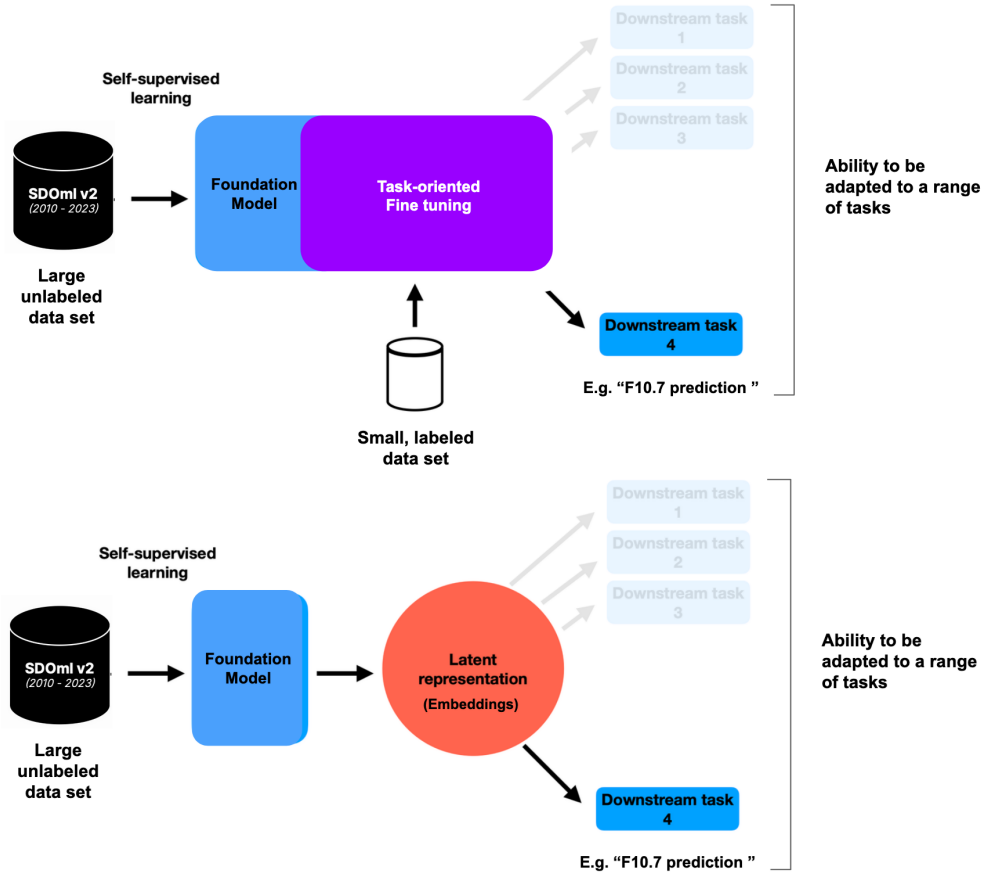


Figure 1: Two methods of using the pre-trained backbone, directly with an adaptor for fine-tuning [10], or with a new model consuming the generated latent representation directly.

3 Method

Our foundation model is composed of a *backbone*, optional *neck*, and *head*. We define the backbone as the model initially trained on the reconstruction task, the neck as the converter between backbone and head, the former then selected for the downstream application (or validation task). We implement two model families as backbones, one stemming from a Nouveau Variational Autoencoder (NVAE) [11], the other from a MAE [12]. They are both adapted to better accommodate our scientific dataset and for intermediate export of their latent spaces in the form of “embeddings.” We additionally evaluated various feature engineering options regarding how to manage the solar disk, the most effective included a simple look up from Stonyhurst coordinates, a heliographic coordinate system for a fixed observer on Earth (suitable given the geosynchronous orbit), to pixel space.

3.1 Model choice

Model selection was initially determined by the ability to capture solar phenomena, guided by applicability to SDO imagery and the ease of access to the embeddings in the latent space. The autoencoder architecture was selected for the backbone for ease of embedding construction and extraction. By design, autoencoders create a lower-dimensional representation during the encoding process. Other requirements included engineering efficiencies, such as ability to mask the solar limb for on-disk experiments, and cheaply bias by importance sampling for areas of interest (e.g. active regions).




Figure 2: Visualizations of the SAMAE (left) and Nouveau-VAE (right), where the SAMAE input has been transformed for full coverage, and the original Nouveau-VAE code base expanded to enable extraction of the latent representation.

3.1.1 Solar-aware Masked Autoencoder

Masked Autoencoders (MAEs) learn to be capable at reconstructing images with random components removed [13]. The approach follows the standard ViT-patchification common to transformer computer vision approaches for deconstruction of the image that the attention mechanism can learn between. The source of this “powerful expressivity” is attributed to “rich hidden representation” [14]. This is particularly of interest in our scenario, as we seek to learn which components of solar imagery are of value for our scientific validation cases. This model was expanded to increase suitability for temporal information for remote sensing tasks [15]. We have continued to iterate, adding “solar-awareness” by including the ability to process the nine wavelengths of interest to us via the Atmospheric Imaging Assembly, efficiencies for processing the solar disk, and the ability to optionally bias the model towards learning active regions of scientific interest.

3.1.2 Nouveau-VAE

The Nouveau Variational Autoencoder (NVAE) is a deep hierarchical VAE created for image generation. Like the MAE, it is able to create a rich latent space using depth-wise separable convolutions and batch normalization. The NVIDIA team’s codebase was modified to permit access to the hierarchical structure to successfully extract embeddings fig. 2.

3.2 Scientific Validation Cases

Predict F10.7 This index is a proxy for solar irradiance, which can be measured from the ground, as this frequency is not absorbed by the atmosphere. Can we achieve good agreement with ground measurements? There is limited scientific value in this prediction of a proxy measure such as F10.7, however this simple task clearly indicates learned capacity in a single result.

Virtual EVE In 2014, an instrument malfunction resulted in the loss of the MEGS-A module of SDO/EVE. With four years of overlapping data, [16, 17] used a hybrid CNN/linear regression model to successfully demonstrate the capability of machine learning methods to estimate missing EUV irradiance measurements from MEGS-A (and the degraded MEGS-B components of the EVE instrument). This validation task employs the embeddings constructed from AIA to understand the contributions from solar features on the EUV spectra, as the mapping between instruments exists due to the narrow-band images (SDO/AIA) and sun-as-a-star spectra (SDO/EVE) observing the same plasma distribution. A linear model accounts for a large portion of the relationship, while a CNN is used to correct for outlier events such as solar flares. There are known concerns regarding the model’s performance post-2020, as AIA instrument performance deviates further from the 2014 baseline. Some of these issues can be addressed by incorporating other sources of irradiance, such as data

from sounding rockets, for training over longer periods, although these are sparse. Importantly, this outperforms a physics-based inversion approach [18].

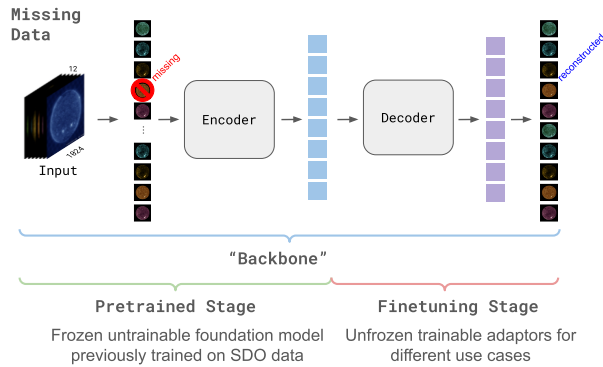


Figure 3: Missing/corrupt data reconstruction process.

Missing Channel Reconstruction The reconstruction of missing extreme ultraviolet (EUV) images from wavelength images is a crucial task given the often low or unusable quality of image data frames from the Solar Dynamics Observatory (SDO). Currently, there is no effective method to recover these missing steps. However, the foundation model is capable of reconstructing individual frames by leveraging contextual information available in other wavelength channels. This approach allows for interpolation to provide a best-guess estimate of missing data at any arbitrary time step.

As with the Virtual EVE project, and differential emission measure analysis [18], the overlapping temperature range covered by different SDO/AIA wavelength channels allows for the temperature distribution of the underlying plasma to be reconstructed, may enable the inference of properties of different temperature ranges.

The overlapping range covered by different wavelength channels may enable the inference of properties of different temperature ranges. This overlap can be used within a machine learning model to produce an estimation to replace data that is either missing, corrupted, or otherwise unusable. Our objective is to develop a more robust model that operates with higher computational efficiency while producing results comparable to the current SOTA. Special attention is given to the model's ability to capture non-linear relationships or rare events, such as intensity values in flaring regions.

There are several uncertainties inherent in this process. Some channels may be more readily recreated than others due to the physical assumptions that channels in the middle of the temperature/wavelength ranges will have the most overlap with other channels, potentially yielding better results. However, this overlap might not always correspond to the actual missing data in the SDO. Addressing these uncertainties requires an understanding of the shortfalls to determine the appropriateness of this reconstruction technique in different scenarios.

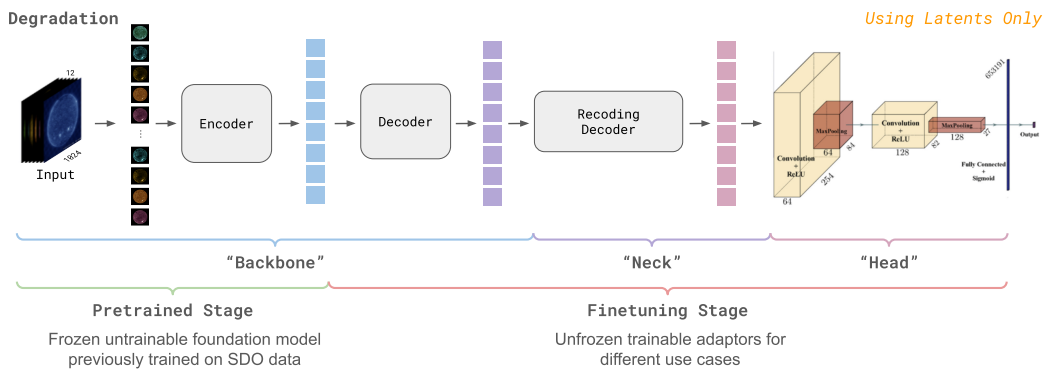


Figure 4: Instrument degradation prediction, head architecture reproduced with permission, [19].

Autocalibration The SDO/AIA EUV channels exhibit degradation due to exposure to the same emissions they are intended to measure. This degradation results in apparent dimming over time across multiple EUV channels with unique characteristics. This poses challenges for long-term studies, as degradation trends within the dataset need to be corrected. Until 2014, SDO utilized EVE to correct this degradation. As discussed, a malfunction of SDO/EVE resulted in the loss of the MEGS-A component, and calibration is currently performed by sounding rocket flights. In response to this, [19] used a CNN to reconstruct the Atmospheric Imaging Assembly (AIA) multi-channel degradation curves.

Data requirements for this study include the SDOML data from AIA as well as older correction tables. The sampling requirement is minimal, with data being required once per day or even less frequently. Traditional SOTA methods, such as those performed by the Lockheed Martin Solar and Astrophysics Laboratory (LMSAL), involve calibration using sounding rocket flights. These methods, while accurate, are expensive and technically demanding. Our goal is to reproduce the results from [19] with greater efficiency in terms of data required and computational resources. This efficiency is evaluated through an examination of the resultant images compared to those produced by SOTA calibration pipelines, alongside intensity histograms, data spike analysis, and other metrics.

4 Results

Overall, our model families were evaluated for their backbone reconstruction task and against our four scientific validation cases. In all but the autocalibration task they reached the same level of accuracy or surpassed their classical counterparts in a fraction of the required time. In the autocalibration case, the direct embedding approach was able to match, but took additional training time.

4.1 Reconstruction

Loss for the reconstruction task is measured by pixel RMSE within the solar disk. SAMAE results presented in fig. 5 indicate a clear ability to reconstruct most wavelengths under a *small* embedding dimension (128) and within a *short* number of training epochs (10). Interestingly this model struggles to reconstruct 131 & 171 Å, which is likely due to a normalization error we're still investigating. The Nouveau-VAE model on raw pixel intensity performs better, even when including the solar limb fig. 6.

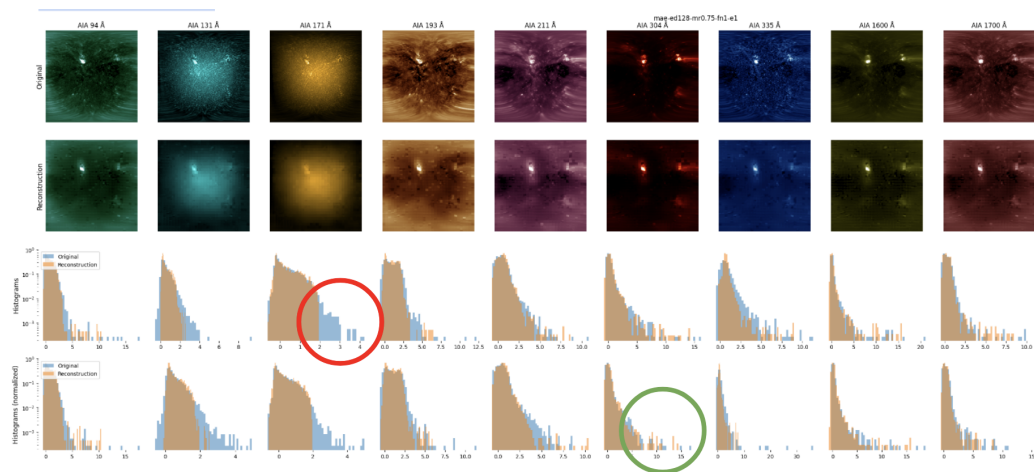


Figure 5: SAMAE reconstruction with disk transform, notice how in the circled areas peaks of some wavelengths are not captured well.

4.2 Direct Embeddings

Training each scientific validation case on the embeddings directly led to generally much faster training time and matching or surpassing of accuracy. There was an effort made to best evaluate

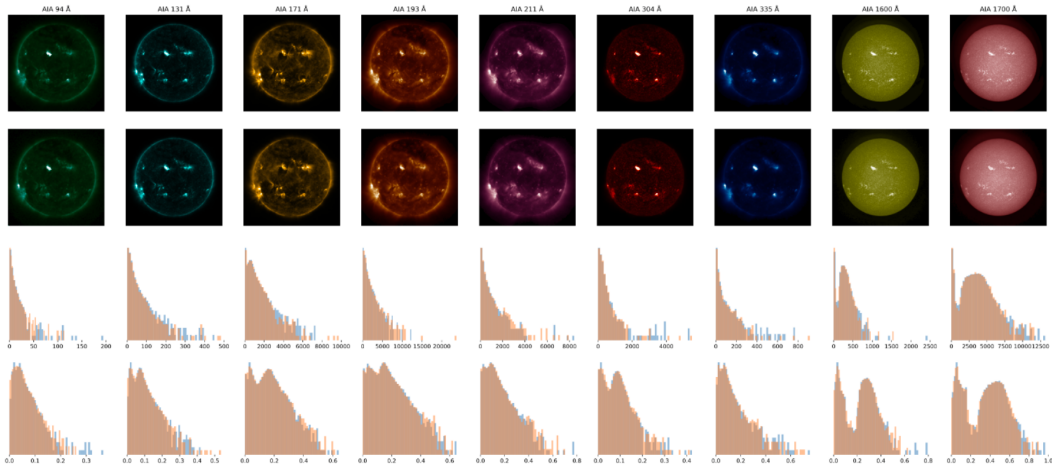


Figure 6: Nouvaeu-VAE reconstruction, with better performance.

the embeddings outside of the scientific cases to consider embedding-to-embedding comparison. The common TSNE approach was taken over a small one-year sample, fig. 9 indicates there was separation by solar activity. This approach however is still fairly opaque and hence the validation approaches are considered more appropriate.

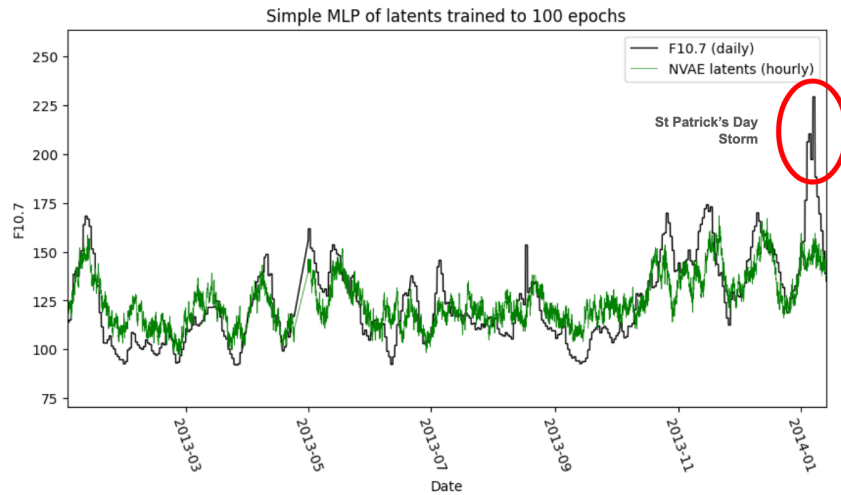


Figure 7: Direct embedding prediction of F10.7 proxy.

Scaling Scaling engineering effort on TPUs is a different set of skills than ML-engineering effort on GPUs. On GPUs the engineering allows for rapid model development and innovation; often possible on a local GPU machine. In contrast, TPUs are less accessible; requiring significant engineering to build a basic framework and often need to be reserved in advance through Google Cloud. The decision between GPU and TPU would be easier if there was mature, hardware agnostic layers. Tools like PyTorch Lightning show promise in this area but still require significant engineering to develop.

Finally, model training can be extrapolated roughly from historical training times. If we take a most closed-form scenario whereby one epoch takes 90 minutes to process 3 months of SDO AIA on a TPU with $16 \times 8 = 128\text{Gb}$ available memory at \$1.5 per chip hour, scaling over various data training periods would look roughly like fig. 8 and require on-disk requirements and VRAM to accommodate estimated in table 1.

Training Time vs TPU Size

3 data scales; 100 epochs




Figure 8: TPU scaling training time estimations over different sizes of TPU and total time period.

Parameter Size	Precision		Size on disk (Gb)		Size on GPU (Gb)		
	<i>BF16</i>	32	<i>BF16</i>	32	<i>Quantized 8-bit</i>	<i>BF16</i>	32
<i>330M</i>	5B	11B	0.66	1.32	0.36	0.73	1.45
<i>700M</i>	11B	22B	1.40	2.80	0.77	1.54	3.08
<i>2B</i>	32B	64B	4.00	8.00	2.20	4.40	8.80
<i>10B</i>	160B	320B	20.00	40.00	11.00	22.00	44.00

Table 1: Size-on-disk and VRAM requirement estimates by increasing parameter size.

5 Discussion

The adaptation of foundation models (FMs) to diverse scientific use-cases is a crucial step toward unlocking the full potential of data-driven research in space sciences. This initial exploration with the Solar Dynamics Observatory Foundation Model (SDO-FM) showcases the flexibility of these models within the context of space science, demonstrating that this new generation of large scale models can be successfully tailored to specific downstream objectives through two primary strategies: fine-tuning and leveraging latent representations. This research highlights the potential for reconfiguring embeddings to address a wide range of heliophysics tasks without the need for exhaustive retraining, simplifying the cost and complexity of using ML for heliophysics and space weather investigations.

The first strategy: selective fine-tuning, involves selectively freezing portions of the backbone model while allowing specific layers, often including the task-specific head, to continue learning. This selective fine-tuning allows for the preservation of learned representations in the backbone, while enabling specialization to new downstream tasks. By maintaining the core structure of the FM, we reduce the risk of catastrophic forgetting, a frequent challenge in transfer learning (where new learned representations are updated and optimized for the new task at the expense of prior knowledge) and optimize for both efficiency and precision. This approach is especially beneficial for scientific use-cases where domain-specific generalizations can be complex and retraining an entire model from scratch would be computationally prohibitive.




Figure 9: TSNE applied to a small embedding sample, colored top left supporting to higher activity months.

The second strategy: leveraging latent representations, offers a computationally efficient approach by bypassing the need for traditional fine-tuning. Instead, we extract compressed latent embeddings from the backbone, which capture high-level, abstract features of the input data. These latent embeddings can then serve as a basis for new models focused on downstream tasks such as reconstruction or prediction. This approach also drastically reduces the size of the SDO dataset, simplifying the input to subsequent models and handling complexity while maintaining the essential phenomenological information. Latent feature extraction is particularly valuable in environments where computational resources are limited; for example on a spacecraft; as it allows for faster training without adversely compromising outcomes.

Both the selective fine-tuning and latent representation strategies were effective in adapting SDO-FM to downstream scientific tasks, marking a significant advancement in the development of foundation models using multi-modal space-acquired data. Our findings underscore the importance of structured governance in the ongoing development and deployment of foundation models for science, of which SDO-FM is an early exemplar. Foundation models are a new class of scientific tool and, in line with the principles guiding traditional scientific instruments, we propose that all science-based foundation models should be overseen by a principal investigator (PI), supported by a multidisciplinary working group. This group would include scientists, data providers, and data scientists, ensuring broad awareness of the model's capabilities, effective onboarding, and clear communication regarding versioning and expectations. For instance, this initial version of SDO-FM, with its relatively low temporal resolution can only emulate proven scientific results. Thus, it is essential for future versions to have clearly defined science goals and user expectations, where the science requirements are driving the engineering choices.

Adherence to open science and transparency around model uncertainty should be central to continued development of foundation models. Comprehensive validation protocols, such as those established by Munoz-Jaramillo's community-built framework for forecasting should be integral to the model's development, analogous to aerospace industry flight checks, particularly in cases where there may be imbalances in the data. Moreover, benchmarking SDO-FM against other foundation models and assessing its performance across both in-distribution and out-of-distribution scenarios will be critical for ensuring its scientific robustness. Shared definitions of progress, coupled with transparent publication of failure modes in accordance with NASA's Transform to Open Science initiative, will facilitate ongoing improvements, accountability, and trust within the scientific community and ensure this shared resource is used. Ultimately, the development of a robust, adaptable foundation model in the field of space science hinges on a coordinated effort that integrates scientific rigor, engineering best practices, and an unwavering commitment to transparency and accessibility.

Acknowledgments and Disclosure of Funding

This work is the research product of the SDO-FM: A Multi-Modal Foundation Model POC for SDO (Grant#: 80NSSC24K0701); funded and supported by NASA. The research and its outputs have been designed, managed, and delivered by Trillium Technologies Inc.

References

- [1] P. H. Scherrer, J. Schou, R. I. Bush, A. G. Kosovichev, R. S. Bogart, J. T. Hoeksema, Y. Liu, T. L. Duvall, J. Zhao, A. M. Title, C. J. Schrijver, T. D. Tarbell, and S. Tomczyk. The Helioseismic and Magnetic Imager (HMI) Investigation for the Solar Dynamics Observatory (SDO). *Solar Physics*, 275(1-2):207–227, January 2012.
- [2] James R. Lemen, Alan M. Title, David J. Akin, Paul F. Boerner, Catherine Chou, Jerry F. Drake, Dexter W. Duncan, Christopher G. Edwards, Frank M. Friedlaender, Gary F. Heyman, Neal E. Hurlburt, Noah L. Katz, Gary D. Kushner, Michael Levay, Russell W. Lindgren, Dnyanesh P. Mathur, Edward L. McFeaters, Sarah Mitchell, Roger A. Rehse, Carolus J. Schrijver, Larry A. Springer, Robert A. Stern, Theodore D. Tarbell, Jean-Pierre Wuelser, C. Jacob Wolfson, Carl Yanari, Jay A. Bookbinder, Peter N. Cheimets, David Caldwell, Edward E. Deluca, Richard Gates, Leon Golub, Sang Park, William A. Podgorski, Rock I. Bush, Philip H. Scherrer, Mark A. Gummin, Peter Smith, Gary Aufer, Paul Jerram, Peter Pool, Regina Souffi, David L. Windt, Sarah Beardsley, Matthew Clapp, James Lang, and Nicholas Waltham. The Atmospheric

- Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO). *Solar Physics*, 275(1-2):17–40, January 2012.
- [3] T. N. Woods, F. G. Eparvier, R. Hock, A. R. Jones, D. Woodraska, D. Judge, L. Didkovsky, J. Lean, J. Mariska, H. Warren, D. McMullin, P. Chamberlin, G. Berthiaume, S. Bailey, T. Fuller-Rowell, J. Sojka, W. K. Tobiska, and R. Viereck. Extreme Ultraviolet Variability Experiment (EVE) on the Solar Dynamics Observatory (SDO): Overview of Science Objectives, Instrument Design, Data Products, and Model Developments. *Solar Physics*, 275(1-2):115–143, January 2012.
 - [4] Richard Galvez, David F Fouhey, Meng Jin, Alexandre Szenicer, Andrés Muñoz-Jaramillo, Mark CM Cheung, Paul J Wright, Monica G Bobra, Yang Liu, James Mason, et al. A machine-learning data set prepared from the nasa solar dynamics observatory mission. *The Astrophysical Journal Supplement Series*, 242(1):7, 2019.
 - [5] Will T. Barnes, Mark C. m. Cheung, Monica G. Bobra, Paul F. Boerner, Georgios Chintzoglou, Drew Leonard, Stuart J. Mumford, Nicholas Padmanabhan, Albert Y. Shih, Nina Shirman, David Stansby, and Paul J. Wright. aiapy: A python package for analyzing solar euv image data from aia. *Journal of Open Source Software*, 5(55):2801, 2020.
 - [6] Edward Brown, Stefano Bonasera, Bernard Benson, J Pérez-Hernández, Giacomo Acciarini, A Baydin, Christopher Bridges, Meng Jin, Eric Sutton, and Moriba K Jah. Learning the solar latent space: sigma-variational autoencoders for multiple channel solar imaging. In *Fourth workshop on machine learning and the physical sciences (neurips 2021)*, 2021.
 - [7] Shreshth A Malik, James Walsh, Giacomo Acciarini, Thomas E Berger, and Atılım Güneş Baydin. High-cadence thermospheric density estimation enabled by machine learning on solar imagery. *arXiv preprint arXiv:2312.06845*, 2023.
 - [8] Christoph Feichtenhofer, Yanghao Li, Kaiming He, et al. Masked autoencoders as spatiotemporal learners. *Advances in neural information processing systems*, 35:35946–35958, 2022.
 - [9] Johannes Jakubik, Sujit Roy, C. E. Phillips, Paolo Fraccaro, Denys Godwin, Bianca Zadrozny, Daniela Szwarcman, Carlos Gomes, Gabby Nyirjesy, Blair Edwards, Daiki Kimura, Naomi Simumba, Linsong Chu, S. Karthik Mukkavilli, Devyani Lambhate, Kamal Das, Ranjini Bangalore, Dario Oliveira, Michal Muszynski, Kumar Ankur, Muthukumaran Ramasubramanian, Iksha Gurung, Sam Khallaghi, Hanxi, Li, Michael Cecil, Maryam Ahmadi, Fatemeh Kordi, Hamed Alemohammad, Manil Maskey, Raghu Ganti, Kommy Weldemariam, and Rahul Ramachandran. Foundation models for generalist geospatial artificial intelligence, 2023.
 - [10] Anna Scaife. Foundation models for radio astronomy. 2nd ICML Workshop on Machine Learning for Astrophysics, 2024.
 - [11] Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33:19667–19679, 2020.
 - [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.
 - [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021.
 - [14] Shuhao Cao, Peng Xu, and David A. Clifton. How to understand masked autoencoders, 2022.
 - [15] Johannes Jakubik, Sujit Roy, C. E. Phillips, Paolo Fraccaro, Denys Godwin, Bianca Zadrozny, Daniela Szwarcman, Carlos Gomes, Gabby Nyirjesy, Blair Edwards, Daiki Kimura, Naomi Simumba, Linsong Chu, S. Karthik Mukkavilli, Devyani Lambhate, Kamal Das, Ranjini Bangalore, Dario Oliveira, Michal Muszynski, Kumar Ankur, Muthukumaran Ramasubramanian, Iksha Gurung, Sam Khallaghi, Hanxi (Steve) Li, Michael Cecil, Maryam Ahmadi, Fatemeh Kordi, Hamed Alemohammad, Manil Maskey, Raghu Ganti, Kommy Weldemariam, and Rahul Ramachandran. Foundation Models for Generalist Geospatial Artificial Intelligence. *Preprint Available on arxiv:2310.18660*, October 2023.

- [16] Alexandre Szenicer, David F. Fouhey, Andres Munoz-Jaramillo, Paul J. Wright, Rajat Thomas, Richard Galvez, Meng Jin, and Mark C. M. Cheung. A deep learning virtual instrument for monitoring extreme UV solar spectral irradiance. *Science Advances*, 5(10):eaaw6548, October 2019.
- [17] Manuel Indaco, Daniel Gass, William James Fawcett, Richard Galvez, Paul J. Wright, and Andrés Muñoz-Jaramillo. Virtual eve: a deep learning model for solar irradiance prediction, 2024.
- [18] Paul J. Wright, Mark C. M. Cheung, Rajat Thomas, Richard Galvez, Alexandre Szenicer, Meng Jin, Andrés Muñoz-Jaramillo, and David Fouhey. DeepEM: Demonstrating a Deep Learning Approach to DEM Inversion, March 2019.
- [19] Luiz F. G. Dos Santos, Souvik Bose, Valentina Salvatelli, Brad Neuberg, Mark C. M. Cheung, Miho Janvier, Meng Jin, Yarin Gal, Paul Boerner, and Atılım Güneş Baydin. Multichannel autocalibration for the Atmospheric Imaging Assembly using machine learning. *Astronomy & Astrophysics*, 648:A53, April 2021.