



A Machine-learning Data Set Prepared from the NASA *Solar Dynamics Observatory* Mission

Richard Galvez¹ , David F. Fouhey² , Meng Jin^{3,4} , Alexandre Szenicer⁵ , Andrés Muñoz-Jaramillo⁶ , Mark C. M. Cheung^{3,7} , Paul J. Wright⁸ , Monica G. Bobra⁷ , Yang Liu⁷ , James Mason⁹ , and Rajat Thomas¹⁰

¹Center for Data Science, New York University, New York, NY 10011, USA; richard.galvez@nyu.edu

²University of Michigan, Ann Arbor, MI 48109, USA

³Lockheed Martin Solar & Astrophysics Laboratory, Palo Alto, CA, USA

⁴SETI Institute, Mountain View, CA 94043, USA

⁵University of Oxford, Oxford OX1 2JD, UK

⁶Southwest Research Institute, San Antonio, TX 78238, USA

⁷Hansen Experimental Physics Laboratory, Stanford University, Stanford, CA 94305, USA

⁸SUPA School of Physics & Astronomy, University of Glasgow, Glasgow G12 8QQ, UK

⁹NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

¹⁰University of Amsterdam, 1012 WX Amsterdam, Netherlands

Received 2019 January 19; revised 2019 March 11; accepted 2019 March 13; published 2019 May 8

Abstract

In this paper, we present a curated data set from the NASA *Solar Dynamics Observatory* (*SDO*) mission in a format suitable for machine-learning research. Beginning from level 1 scientific products we have processed various instrumental corrections, down-sampled to manageable spatial and temporal resolutions, and synchronized observations spatially and temporally. We illustrate the use of this data set with two example applications: forecasting future extreme ultraviolet (EUV) Variability Experiment (EVE) irradiance from present EVE irradiance and translating Helioseismic and Magnetic Imager observations into Atmospheric Imaging Assembly observations. For each application, we provide metrics and baselines for future model comparison. We anticipate this curated data set will facilitate machine-learning research in heliophysics and the physical sciences generally, increasing the scientific return of the *SDO* mission. This work is a direct result of the 2018 NASA Frontier Development Laboratory Program. Please see the Appendix for access to the data set, totaling 6.5TBs.

Key words: astronomical databases: miscellaneous – catalogs – editorials, notices – miscellaneous – surveys

1. Introduction

Launched in 2010, NASA's *Solar Dynamics Observatory* (*SDO*; Pesnell et al. 2012) has been continuously monitoring the Sun's activity and delivering valuable scientific data for heliophysics researchers with the use of three instruments:

1. The Atmospheric Imaging Assembly (AIA; Lemen et al. 2012), which captures 4096×4096 resolution images (with 0.6 arcsec pixel size) of the full Sun in two ultraviolet (UV; centered at 1600 and 1700 Å) wavelength bands, seven extreme ultraviolet (EUV) wavelength bands (centered at 94, 131, 171, 193, 211, 304, and 335 Å) and one visible wavelength (centered at 4500 Å).
2. The Helioseismic and Magnetic Imager (HMI; Schou et al. 2012) captures visible wavelength filtergrams of the full Sun at 4096×4096 resolution (a pixel size of 0.5 arcsec), which are then processed into a number of products, including photospheric Dopplergrams, line-of-sight magnetograms, and vector magnetograms (Hoeksema et al. 2014).
3. The EUV Variability Experiment (EVE; Woods et al. 2012) monitors the solar EUV spectral irradiance from 1 to 1050 Å. This is done by utilizing multiple EUV Grating Spectrographs (MEGS) that disperse EUV light

from the full disk of the Sun and its corona onto a 1024×2048 charge coupled device (CCD).

Calibrated level 1 scientific data from the AIA and HMI instruments are accessible from the Joint Science Operations Center¹¹ (JSOC) at Stanford University, Lockheed Martin Solar & Astrophysics Laboratory, and affiliate science data centers, while science data from the EVE instrument are accessible from the EVE Science Operations Center¹² at the Laboratory for Atmospheric and Space Physics at the University of Colorado, Boulder.

The *SDO* mission has been scientifically prolific. In the eight years after launch, over 3000 refereed scientific publications¹³ have made use of *SDO* data. This success can be attributed to the reliability of the spacecraft and its instruments, the consistency and quality of the observations, the mission's open data policy, and the ease of online data access from the affiliated science data centers. The large volume of structured, calibrated scientific data (over 12 Petabytes and counting) is poised for an exploratory analysis from machine-learning methods, as well as more traditional approaches. In early pioneering works, supervised learning techniques have been applied to the prediction of solar flares using HMI vector magnetograms (e.g., Bobra & Couvidat 2015), as well as HMI and AIA imagery in Jonas et al. (2018). Deep-learning applications have begun to emerge from the heliophysics



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

¹¹ <http://jsoc.stanford.edu>

¹² <http://lasp.colorado.edu/home/eve/data>

¹³ <https://sdo.gsfc.nasa.gov/mission/publications.php>

community as well, with exemplary cases illustrated in Colak & Qahwaji (2013) and Huang et al. (2018), with Wright et al. (2019) presenting a more recent treatment using the data set presented here.

While level 1 data are easily accessible, preprocessing these data for a scientific analysis often requires specialized heliophysics knowledge. The necessity for such preprocessing may act as an unnecessary hurdle for non-heliophysics machine-learning researchers who may wish to experiment with data sets from the physical sciences but are unaware of domain-specific nuances (e.g., that images must be spatially and temporally adjusted).

The first contribution of this paper is a curated *SDO* data set that is mission-ready for machine-learning applications. Our aim is to supply this standardized data set for heliophysicists who wish to use machine learning in their own research, as well as machine-learning researchers who wish to develop models specialized for the physical sciences. In Section 2, we examine current available data products and the pitfalls for their direct use in machine-learning tasks, as well as what corrections and adjustments they warrant. These corrections are incorporated into our data preparation procedures discussed in Section 3.

The second contribution of this paper are protocols, metrics, and baseline models. We introduce evaluation protocols and metrics in Section 4 and baseline models in Section 5 where we tackle the tasks of predicting irradiance using present and future EVE data, as well as translating three HMI channels into nine AIA channels. We believe these models contain generic enough components while providing useful benchmarks, and highlighting the most dangerous pitfalls, for most subsequent *SDO* machine-learning applications.

By providing these standardized data products along with accompanying protocols, metrics, and baselines, our aim is to remove unnecessary hurdles for future machine-learning research in heliophysics and the physical sciences more broadly.

2. Examination of Raw Data Products

We first examine existing raw data products available from *SDO* for each of the three instruments (level 1 science data products for AIA, `hmi.B_720s` for HMI, and the EVE version 5 Interactive Data Language (IDL) saveset).

While heliophysics researchers are likely aware of corrections that must be applied to this data, and the fact that AIA measurements have heterogeneous exposure times, it is unrealistic to expect the same from researchers in other fields (e.g., the data set of Kucuk et al. 2017 was compiled from quick-look JPEG2000 images that have a compressed dynamic range and does not account for instrumental degradation). We therefore process these corrections by identifying and removing corrupt observations (e.g., images taken during instrument anomalies), adjusting detected intensities for heterogeneous exposure times, and fixing instrument artifacts that introduce spurious trends.

If such corrupt observations or various sources of heterogeneity are not removed, any subsequent machine-learning model will likely learn to emulate these incorrect observations, as well as any spurious trends, and will not be able to isolate the physical dynamics. Exposing such corrupt data during model training may also compromise predictive quality—or worse, the model may even learn to emulate nonphysical

aberrations and instrumental noise. See Figure 1 for an example of one such unwanted AIA observation.

To identify each instrument’s possible issues, we visualize each instrument’s data by taking the average channel values (i.e., AIA wavelength band data counts, HMI vector field components, and EVE irradiance values) and plot them over time. We then identify the underlying causes of nonphysical aberrations and what necessary corrections are needed to standardize the data. Below, we report our analysis for AIA and outline where HMI required similar adjustments; EVE level 3 data products already address all instrumental issues so we only adjusted for time synchronicity. We describe these corrections in Section 3.

The average channel values for the AIA level 1 data products, as plotted in Figure 2, shows the data heterogeneity as well as the presence of corrupt observations. These are visible in this figure as isolated downward spikes, while the secular downward trend is indicative of degradation over the lifetime of the instrument.

Corrupt observations arise due to a variety of reasons, such as data reported during calibration maneuvers, eclipse periods, or the occasional instrument anomaly. Such data, flagged by a nonzero value of the `QUALITY` keyword for both the AIA and HMI instruments, are not intended for a scientific analysis and are removed from our data set. One of the main sources of heterogeneity in AIA data responds to its instrument design: AIA instrumentation is not designed to directly measure irradiance, but rather data numbers (DNs) as tabulated by the activation of the CCD instrument. While it is intended that DN values are proportional to the flux of photons at a specific wavelength (Boerner et al. 2012), the factor of proportionality is not constant in time. For instance, the camera exposure time, t_{exp} , is not constant due to instrumental automatic exposure control (AEC); e.g., in times of flares when certain regions on the Sun become especially bright, AEC reduces the nominal exposure time from a few seconds to tens of milliseconds. Due to these factors, when the AEC is activated, the mean registered DNs are drastically reduced, which is an effect easily compensated for by adjusting for the exposure time.

The visible downward trend in Figure 2 is caused by the gradual in-orbit degradation of the AIA instrument. This degradation is purely due to CCD corrosion over time. Because AIA is calibrated against EVE, which is itself bootstrapped to a program of regular EUV spectral irradiances as measured from sounding rocket launches (Boerner et al. 2014), the time-dependent profile of the AIA instrument is well understood independently of the solar cycle. This instrumental understanding allows us to correct for AIA instrument degradation by simply applying the `aia_get_response` routine in the *SolarSoft* software package (Freeland & Handy 1998).

Lastly, there is a more subtle nonmonotonic heterogeneity caused by *SDO*’s orbit around the Sun. *SDO* is in a geosynchronous orbit around the Earth, which itself is in a slightly elliptical orbit around the solar system’s barycenter. The elliptical orbit causes the size of the Sun (in DNs registered on the CCD) to vary by about 10% over the course of a year. This is not an intrinsic feature of solar evolution. We compensate for this effect by resizing AIA and HMI images such that the size of the solar disk is fixed to some size R_s . In particular, we can scale the AIA and HMI images by a factor of

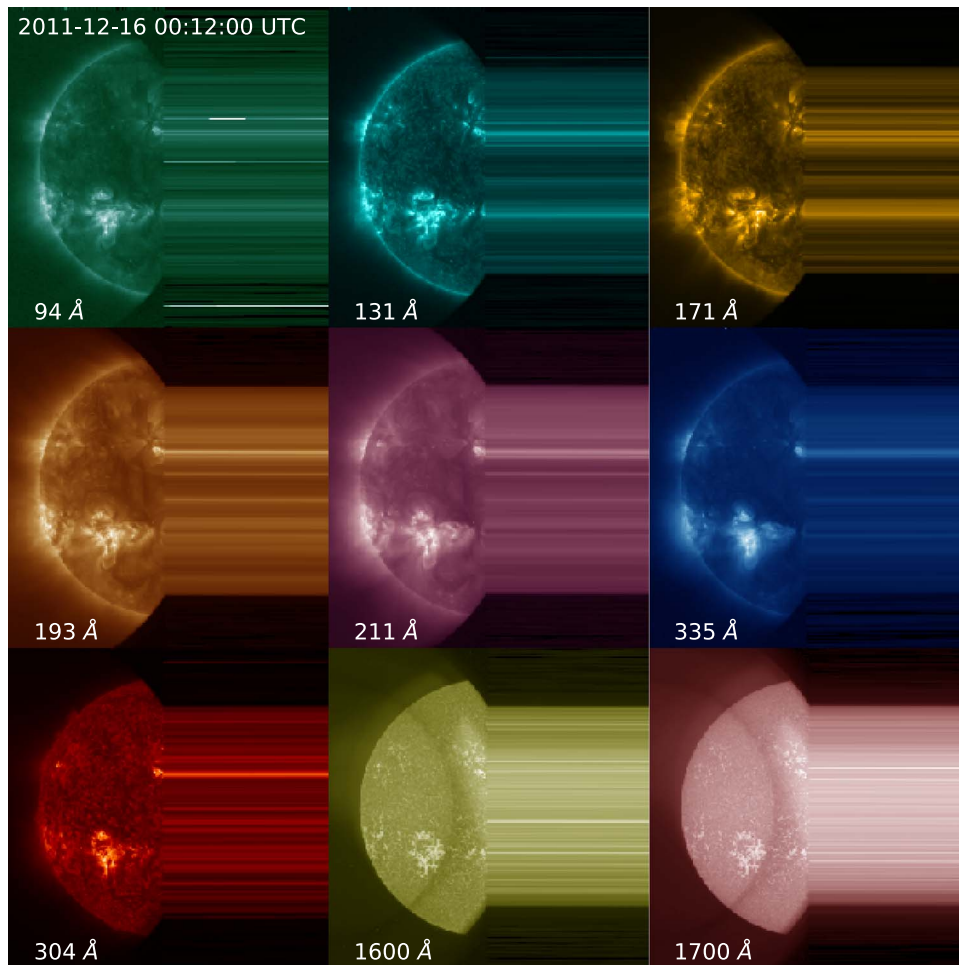


Figure 1. Example of a corrupt observation from the AIA instrument. Utilizing such observations during the training phase of a machine-learning model may compromise its predictive capability.




Figure 2. Average AIA wavelength band data counts for level 1 data products over time, presented here as randomly down-sampled to 1000 observations. The secular downward trend is caused by instrument degradation over time, while the spurious per-channel drops are caused by the instrument’s automatic exposure control mechanism.

R_s/R_{Obs} , where R_{Obs} can be obtained from the RSUN_OBS keyword in the level 1 FITS header.

3. Processed Data Preparation

We now describe how our processed data set is produced in detail. First, we describe how the needed corrections outlined in Section 2 are applied to each instrument and how temporal synchronization is computed. We first begin by removing the nonzero QUALITY observations from both AIA and HMI. We then spatially down-sample to produce a more manageable data set while being careful to emulate what lower-resolution instruments would observe.

3.1. AIA

We begin with the 4096×4096 pixel level 1 (Lemen et al. 2012; dark-subtracted, flat-fielded, and despiked) data products and process them as described below:

1. The raw images are rotated and resized onto a common grid (still 4096×4096 pixels) such that the pixel size is 0.600 arcsec, and \hat{x} and \hat{y} (the first and second image dimensions) are aligned with the solar west and north directions, respectively.
2. Images are rebinned by averaging neighboring 4×4 pixel blocks such that the resultant image has a size of 1024×1024 pixels (with a final pixel size of 2.400

arcsec). Resultant images are processed at a 2-minute cadence, producing the so-called Synoptic series.¹⁴

3. The AIA images are then normalized by exposure time and corrected for instrument degradation, while corrections for elliptical orbital variation are applied with a fixed disk size of R_s of 976 arcsec.
4. Finally, the images are down-sampled again by summing in local blocks, which emulates the expected observation of a lower-resolution instrument. The final interpolated images have 512×512 pixels with pixel size of ~ 4.8 arcsec.

3.2. HMI

We start with the original HMI JSOC data series `hmi.B_720s`, which provides the magnetic vector field strength, inclination angle, and azimuth (Hoeksema et al. 2014). We process this to calculate full-disk vector field observations in B_x , B_y , and B_z components with 12-minute cadence. The $+x$ direction points to the solar west, $+y$ to the north, and $+z$ out of the image plane (i.e., line-of-sight). Additionally, as with the AIA instrument, although the original image size is 4096×4096 , the pixel resolution is different. We therefore further co-aligned HMI and AIA data so that they have the same spatial sampling. The major processing steps for the HMI observations are as follows:

1. We begin by converting the original HMI JSOC data series `hmi.B_720s` vector field data with the disambiguation solution of `disambig.fits` to B_x , B_y , and B_z components, spatially co-aligning with AIA observations using the FITS header information.
2. The HMI images were also corrected for orbital variation with a fixed disk size R_s of 976 arcsec throughout.
3. Finally, we down-sampled the data by averaging in local blocks, which emulates the expected observation at the target lower-resolution. The final interpolated images have 512×512 pixels with a pixel size of ~ 4.8 arcsec.

3.3. EVE

EVE spectra are assembled from a battery of instruments, including the MEGS-A, -B, -P, Solar Aspect Monitor, and the EUV SpectroPhotometer. Each of these instruments covers a different wavelength range in the EUV spectra and are cross-calibrated to produce EVE's data products.

The EVE data released in this data set are extracted from a specially prepared EVE version 5 IDL saveset, including 39 emission lines (see Table 1) during the time window between 2010 May 1 and 2014 May 26. The end date of this data set corresponds to the failure of the MEGS-A instrument, which covered the range between 30 and 370 Å. The EVE data have already been calibrated with physical units of Wm^{-2} , scaled to 1 au, and corrected for degradation, requiring no subsequent calibration. The only processing we perform is to convert from IDL to NumPy Arrays and to temporally synchronize with the AIA and HMI observations.

¹⁴ <http://jsoc.stanford.edu/data/aia/synoptic/>

Table 1
EVE Emission Lines, Their Wavelength, and Temperature of the Emission Plasma

Emission Line	Wavelength	Temperature
Fe XVIII	93.9 Å	6.46×10^6 K
Ne VII	127.7 Å	5.01×10^5 K
Fe VIII	131.2 Å	3.71×10^5 K
Fe XX	132.8 Å	9.33×10^6 K
Ne V	148.7 Å	3.16×10^5 K
Fe IX	171.0 Å	6.46×10^5 K
Fe X	177.2 Å	9.77×10^5 K
Fe XI	180.4 Å	1.17×10^6 K
Fe XII	195.1 Å	1.35×10^6 K
Fe XIII	202.0 Å	1.55×10^6 K
Fe XIV	211.3 Å	1.86×10^6 K
He II	256.3 Å	5.62×10^4 K
Fe XV	284.2 Å	5.01×10^4 K
He II	303.8 Å	1.99×10^6 K
Fe XVI	335.4 Å	2.69×10^6 K
Fe XVI	360.8 Å	2.69×10^6 K
Mg IX	368.1 Å	9.77×10^5 K
Mg IX	443.7 Å	1.00×10^6 K
Ne VII	465.2 Å	3.98×10^5 K
Si XII	499.4 Å	1.99×10^6 K
O III	525.8 Å	7.94×10^4 K
O IV	554.4 Å	1.99×10^5 K
He I	584.0 Å	1.99×10^4 K
O III	599.6 Å	7.94×10^4 K
Mg X	624.9 Å	1.26×10^6 K
O V	629.7 Å	2.51×10^5 K
O II	718.5 Å	6.31×10^4 K
N IV	765.1 Å	1.58×10^5 K
Ne VIII	770.4 Å	6.31×10^5 K
O IV	790.2 Å	1.99×10^5 K
H I	972.5 Å	5.01×10^4 K
C III	977.0 Å	5.01×10^4 K
H I	1025.7 Å	5.01×10^4 K
O VI	1031.9 Å	2.51×10^5 K

Note. The MEGS-A emission lines used in Section 5.1 are indicated in bold.

3.4. Temporal Down-sampling and Synchronization

One of the goals of this paper is to produce a data set that is temporally and spatially synchronized for the three *SDO* data products at manageable resolutions. While our scaling to a fixed solar disk size automatically ensures the spatial synchronization of AIA and HMI, all *SDO* data instruments observe at different cadences (AIA: 2 minutes, HMI: 12 minutes, EVE: 10 s) and are not necessarily aligned in time.

In order to perform the temporal synchronization, we down-sample AIA to a 6-minute cadence and match the nearest EVE observation within a mean/max time window of 8.5 s/12 s. This yields a final data set consisting of AIA and EVE observations each at 6-minute cadence with accompanying HMI observations occurring at every other time step¹⁵ with a 12-minute cadence.

3.5. Data

This produces the final data set for the paper totaling ~ 6.5 TB, made available through the Stanford Digital Repository¹⁶ (please see the Appendix for a list of URLs).

¹⁵ With the exception of nonzero QUALITY observations for either AIA or HMI.

¹⁶ <https://purl.stanford.edu/nk828sc2920>

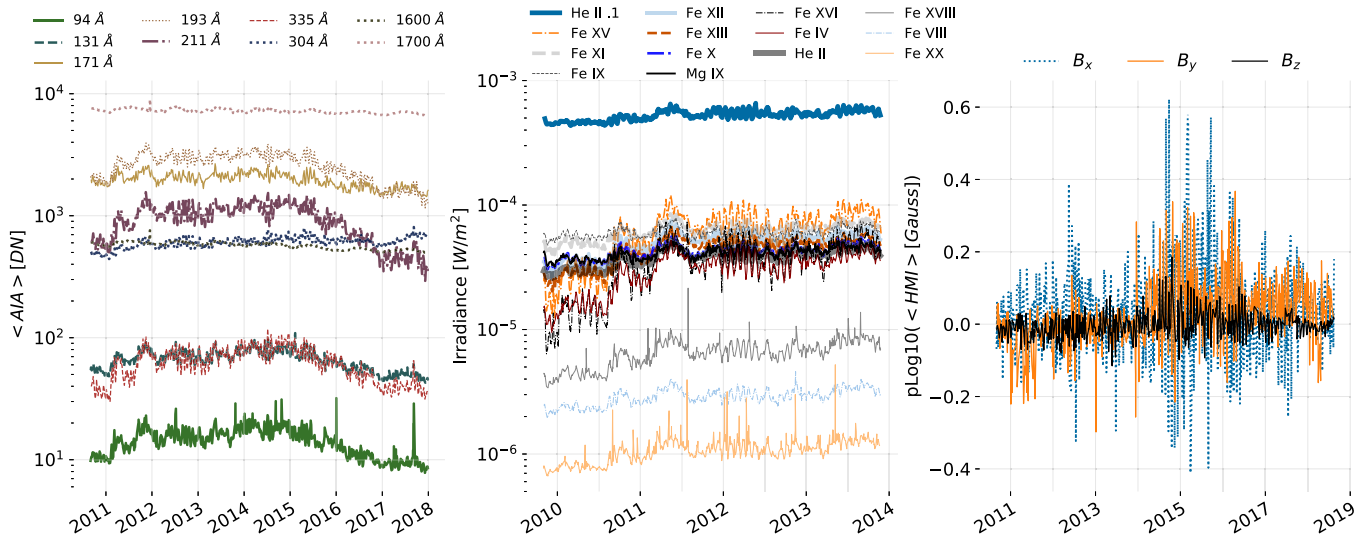


Figure 3. Mean intensity variation after correcting for exposure time, degradation, and orbital variation for AIA (left panel) and EVE (middle panel). Middle panel: we display only the 14 emission lines covered by the MEGS-A instrument, for illustration purposes, out of the 37 in our data product (see Table 1). Right panel: the signed pseudo-logarithm of the mean field values for B_x , B_y , and B_z from HMI after correcting for orbital variation.

The data are individually packed monthly and, for each waveband/component of AIA/HMI, all in the NumPy format. The EVE data are packed in a single TAR file. We show the average value as a function of time for the three products in Figure 3, which demonstrates the removal of spurious trends and artifacts. We also show in Figure 4 an example of co-aligned AIA and HMI observations. The top panel shows the observation near the solar maximum of cycle 24 (2014 February 25 00:00 UT), exhibiting several active regions with strong magnetic field magnitudes and associated EUV emission. The bottom panel shows the observation near the solar minimum of cycle 24 (2018 August 10 00:00 UT), displaying only one active region with a comparatively weak magnetic field and EUV brightness.

4. Protocols and Metrics

We expect that this data will be of interest for machine-learning applications in heliophysics and a simple-access data set for the testing of machine-learning models in the physical sciences. To facilitate this, we have defined standard protocols and metrics to aid future work with this data.

Data splits. There is large temporal coherence in the data since large-scale structures on the Sun evolve at timescales beyond days and months. This leads to issues with randomly sampled splits of the data, which are often done in machine-learning settings with uncorrelated data. In particular, randomly sampled splits will lead to training and testing observations that are separated by days or even minutes. While these observations are indeed distinct points in time, they are generated by virtually the same large-scale structures.

In practice, this means that experiments on randomly split data will be unable to identify overfitting and will likely lead to overly optimistic estimates of generalization performance. The specific issue is that when deploying a model, one tests it on large-scale structures and conditions that are different than the training data. However, if the data is split randomly, the model is never actually evaluated on unseen large-scale structures due to temporal coherence. Therefore, there is no indication of whether the model’s performance is due to generalizing well or

if it is simply explained by the model overfitting to the particularities of the limited large-scale structures observed at the training time.

To preclude this, we have split our data in temporal blocks that break this correlation, consisting of (i) a training set used to fit model parameters (e.g., the filter weights of a convolutional neural network), (ii) a validation set used to set model hyperparameters (e.g., the learning rate for training a network), and (iii) a test set used to evaluate out-of-sample model performance.

All of our data splits are performed over the years (2011–2014), the time period for which all three *SDO* instruments (AIA, HMI, and EVE) were active. This time period provides a data set large enough to support the training of modern models that require copious amounts of data. We set aside years 2012 and 2013 for testing purposes, supplying a wide variety of solar conditions. Years 2011 and 2014 are split into training and validation such that 70% of available EVE observations are used for training (until mid-2011 December) and 30% are used for validation. Of course, other splits are possible, especially for problems not relying on EVE observations. We therefore encourage the community to experiment with various data splits, with the cautionary advice that splits should be constructed in temporal blocks as opposed to random sampling.

Metrics. All of the metrics reported in Section 5 are derived from the normalized absolute error, or $|y_i - \hat{y}_i|/y_i$, where \hat{y} is the model prediction and y is the measurement for data point i . For scalar quantities like EVE prediction, which are intrinsically already averaged over the Sun, we report the average normalized absolute error over all samples in the test set.

For images (e.g., AIA prediction) that are not already spatially integrated, we report a number of metrics. First, we report the average normalized absolute error averaging first over each predicted image’s valid pixels and then over the images. In computer vision research, this average has been noted to often poorly characterize the performance on most pixels (Scharstein & Szeliski 2002) since it can be arbitrarily changed by a small number of large errors. Thus, we also report the percentage of good pixels metric, or the fraction of image pixels with a normalized absolute error less than a fixed percentage, t , for $t = 10\%$, 20% , 50% .




Figure 4. Top panel: co-aligned AIA and HMI data set around the solar maximum (2014 February 25 00:00 UT) of solar cycle 24. Three selected AIA wavebands (171, 193, 211 Å) are shown. Bottom panel: co-aligned AIA and HMI data set around the solar minimum (2018 August 10 00:00 UT) of solar cycle 24. The black circle in the HMI magnetogram shows the location of the solar limb.




Figure 5. Results from the 2-hr EVE Prediction experiment. Left panel: persistence model. Middle panel: average assumption model. Right panel: ridge regression model. This prediction exercise is performed for illustration purposes on 14 MEGS-A emission lines (see Table 1). The forecast errors of all intervals, N hours apart and contained in the years 2012 and 2013, are averaged to produce the average error plotted for an N -hour look-forward time in these figures. The average model surprisingly predicts 7 of the 14 EVE lines within a 10% error and does not show much overall variation, while the persistence model achieves this same performance for 10 lines. The ridge regression model often outperforms the persistence model overall but not in all conditions and not by a substantial margin. The linear and persistence models both show periodic trends consistent with one solar rotation. See Section 5.1 for a discussion.

5. Results

In this section, we provide baseline metrics for simple machine-learning applications utilizing the proposed data set, all implemented in the `PyTorch` library (Paszke et al. 2017). These examples were chosen to illustrate what performance metrics should be expected from future models as well as supplying simple examples for typical use cases. To this end, we have selected and evaluated two problems that demonstrate the temporal nature of the data as well as the alignment between the two spatially resolved sensors: (i) predicting future EVE from present EVE and (ii) transforming HMI observations to AIA observations. These generic models may be reapplied to a wide variety of other problems not discussed in this section: for instance, predicting future AIA from current AIA or predicting EVE from AIA.

We stress that our baselines are not intended to be the top-performing solutions but rather a rubric that shows how well a simple data-driven approach would perform. This serves two functions: first, future model implementations that are more complex should out-perform these baselines in the metrics we propose or other such metrics (e.g., focusing only on flaring events); and second, the baselines provide context necessary to properly evaluate a future model’s performance. For instance, while a more complex model may achieve a low error rate, such as 5%, if our baseline already achieved a similar score, the complexity of the new model may not be warranted.

5.1. EVE-to-EVE Prediction

The goal of this task is to predict future EVE observations given current EVE observations at a future time ranging from a few hours up to a full solar rotation. In order to provide statistically sound benchmarks in light of strong solar variability, we calculate the average relative error over all predictions for look-forward times of various fixed sizes. This statistical approach informs to what extent we can predict overall future phenomena for a given look-forward time, as well as accounts for strong solar variability.

There are two main sources of solar variability for this timescale. In shorter timescales, the main source of variability are flares, which increase the EUV radiative output of the Sun by several orders of magnitude in timescales of minutes and hours. The second is solar rotation itself (27 days at the synodic

Carrington rotation rate). Rotation modulates EUV irradiance because active regions (bright in the EUV) have lifespans of 14–55 days and can come in and out of view as the Sun rotates. This active region permanence induces strong temporal correlations at look-forward times greater than 27 days, as illustrated by the periodicity in Figure 5, as the Sun’s “same face” rotates into view. For model evaluation, we choose a total look-forward time of 29 days, a duration long enough to expose the irradiance periodicity.

For our input data, we use the MEGS-A lines listed in Table 1 with the exception of Fe XVI 361 Å, which is the most sparsely measured line with only $\sim 1\%$ of the average number of line measurements.

Baselines. For this problem, we report three simple baselines:

1. *Persistence.* This model assumes that all future observations of the Sun will be identical to its current state. Thus, for any time jump, it predicts that the future EVE observation will be the same as the current EVE observation.
2. *Constant.* This model assumes that the Sun produces a constant EUV irradiance and therefore gives a constant prediction irrespective of the current EVE observation. We set this constant to the training set average per line.
3. *Linear.* This model assumes that future observations are a linear transformation of the current observations plus a constant bias. We fit this model using ridge regression or a linear model with Tikhonov/L2 regularization. In particular, for a given spectral line and look-forward time, if x_i is the current measurement and y_i is the corresponding future average observation, we solve for w such that $\lambda \|w\|_2^2 + \sum_i^n |w^T x_i - y_i|$ is minimized for all instances, i . We set λ per model by a grid search to minimize the average normalized absolute error, doing two-fold cross validation on the training set.

Results. We evaluate the average normalized absolute error for these models for look-forward times ranging from 2 hr to 29 days in steps of 2 hr and report our results in Figure 5. The linear and persistence both show trends corresponding to the solar rotation: their errors peak at approximately half a solar rotation and reduce steadily until a full rotation occurs, thus confirming the strong correlation between observations




Figure 6. Results for HMI-to-AIA translation. The left panel shows the HMI inputs, while the right panel shows the ground-truth AIA (top panel) as well as the predicted AIA from a 3-layer network (middle panel) and an 11-layer network (bottom panel). While the three-layer network performs well, additional layers (i) reduce artifacts (especially in 131 and 171 Å), presumably due to the depth; and (ii) better resolve coronal predictions (especially in 211 and 304 Å), presumably due to the larger receptive field caused by the additional layers.

separated by exact rotations. The average model’s error, on the other hand, is effectively constant; small variations occur because pairs of 1 day jump observations exist from 2012 January 1 up to 2013 December 30, while pairs of 29 day jumps can only be tested up to 2013 December 2.

Collectively, the results underscore the importance of having good baselines via the surprising effectiveness of even trivial models, such as the persistence or average models. For instance, although the average model entirely ignores the current EVE observation, it is able to predict 7 of the 14 EVE lines with less than 10% of the average normalized absolute error; similarly, at a look-forward time of 27 days, 10 lines can still be predicted within a 10% error by the persistence model.

It is true that the linear model frequently improves on the persistence model, especially for high-error lines like Fe XVI and Fe XV, and look-forward times much less than a full rotation. However, for many look-forward times and lines, the trivial persistence model actually outperforms the relatively complex linear model, demonstrating how simple baselines may assist in properly assessing the effectiveness of a machine-learning model.

5.2. HMI-to-AIA Prediction

We now move on to an example that demonstrates how a convolutional deep-learning model may exploit the spatial richness of our data set. In this application, we show how mapping between the HMI and AIA instruments is learned by treating it as an image-to-image translation problem. Such an approach is common in computer vision research, with applications as diverse as labeling each pixel in the scene with a category label (e.g., building; Shelhamer et al. 2016), generating images from sketches (Isola et al. 2017), inferring 3D properties of scenes (Wang et al. 2015), or detecting the pose of humans (Cao et al. 2017).

We have physical reason to expect that mapping between the HMI and AIA instruments exists. While the HMI instrument infers information about the solar magnetic field from the solar photosphere, the AIA instrument measures UV/EUV emission from the solar chromosphere and corona. Since the chromosphere and corona are spatially structured by the presence of strong magnetic fields, UV/EUV emission will typically reflect

information about the magnetic field through its spatial distribution and vice versa. Here, we show how a simple convolutional model can realize the mapping from HMI to AIA.

Baseline. Our baseline is a deep convolutional neural network. This is a function composed of alternating convolutions and nonlinearities that maps one multichannel image (e.g., a three-channel 256×256 image) to another (e.g., a nine-channel 256×256 image). This function can be fit to a data set of inputs (i.e., HMI) and desired outputs (i.e., AIA) via standard optimization procedures. Throughout, we work with 256×256 images.

We adopt a basic approach for our network consisting of a three parts: (i) an initial feature extraction following ResNet (He et al. 2016), consisting of a 7×7 convolution with stride 2 followed by 3×3 max-pooling with stride 2, which expands the receptive fields of subsequent feature maps; (ii) a variable number of 3×3 convolutions with stride 1; (iii) 3×3 convolution yielding a nine-channel prediction followed by $4 \times$ bilinear up-sampling. All intermediate convolutions have 128 filters and are followed by a rectified linear unit (Nair & Hinton 2010) and batch normalization (Ioffe & Szegedy 2015). By varying the number of intermediate convolutions blocks in part ii, we can control both the parameter count and effective receptive field of the network. We report results with 3, 7, and 11 layers (i.e., with 2, 6, and 10 hidden layers, including the initial convolution in part i).

We train the parameters of the network (e.g., filter weights and biases) via back-propagation and mini-batch stochastic gradient descent (SGD) to minimize the mean-squared error of the prediction. Specifically, we use SGD with Nesterov momentum (Sutskever et al. 2013) with momentum 0.99, weight decay 10^{-8} , and batch size 32. We start with a learning rate of 10^{-3} , which we multiply by 0.1 every 5 epochs, and train for 15 epochs. We checkpoint the network at the end of every epoch and take the network with the lowest validation loss. To help the learning, we divide inputs and outputs per-channel by their average over the training set (i.e., the network is trained to predict the AIA 94 Å image divided by the empirical mean of AIA 94 Å over the training set).

Results. We show sample qualitative results in Figure 6 for 3 and 11 layers. Even with a small number of hidden layers, a

Table 2
Results for HMI-to-AIA Prediction

	Mean (Lower Better)										% Pixels < 10% Error (Higher Better)											
	Avg	94	131	171	193	211	304	335	1600	1700	Avg	94	131	171	193	211	304	335	1600	1700		
3 Layer	2.08	0.80	0.98	4.63	5.39	3.70	0.85	0.72	0.73	0.90	15.4	12.0	13.8	12.1	14.7	12.5	15.8	11.1	22.9	23.5		
7 Layer	0.83	0.35	0.38	1.55	1.94	1.27	0.54	0.45	0.46	0.52	18.1	19.5	17.7	14.3	15.6	14.2	15.2	14.2	24.3	27.7		
11 Layer	0.75	0.37	0.40	1.66	1.55	1.06	0.47	0.37	0.42	0.47	20.8	20.3	20.6	16.4	17.6	16.3	18.2	17.4	28.9	31.4		
					% Pixels < 20% Error									% Pixels < 50% Error (Higher Better)								
3 Layer	29.2	23.7	27.0	23.6	28.8	24.9	31.1	22.2	40.5	40.7	58.0	53.9	56.3	49.8	61.6	57.4	65.3	54.4	62.5	60.6		
7 Layer	34.5	38.3	34.7	27.7	30.4	27.9	30.0	28.2	44.1	49.5	68.9	80.0	74.5	58.0	64.0	61.5	65.1	64.2	74.5	78.7		
11 Layer	39.5	39.2	39.6	31.9	34.3	31.9	35.5	34.6	52.2	56.6	75.0	77.5	76.4	65.1	70.6	68.2	73.8	77.2	81.6	84.3		

Note. The top-performing method is indicated in bold.

simple data-driven approach does a good job of getting the general shape and features of the Sun, which suggests that results that get general features can be explained with relatively simple models and that more complex models must provide additional results compared to this. Adding more layers helps reduce artifacts at the edge of the disk in the 131 and 171 Å channels and helps to more accurately resolve the corona in the 211 and 304 Å channels. The shallower network has difficulty accurately resolving the corona because each prediction is made from a small portion of the Sun; thus, it produces a halolike effect around the entire Sun, as opposed to at specific locations on the disk.

Quantitatively, increasing depth produces strong improvements, as seen in Table 2. With a relatively unsophisticated deep network, 75% of the pixels of AIA across all channels can be predicted within a 50% relative error from HMI observations. As seen by the percentage good pixels metrics, 1600 and 1700 Å observations appear to be among the easier to predict and are almost always a few percentage points higher across both network depths and good pixel thresholds. This serves as a good sanity check on the results, since the photospheric and chromospheric brightness features in these two channels are known to be highly correlated to the photospheric distribution of magnetic fields.

6. Conclusion

In this paper, we present a curated, high-quality data set from all three *SDO* instruments primed for machine-learning research. We have preprocessed this data by down-sampling AIA and HMI images from 4096×4096 to 512×512 pixels, removed $QUALITY \neq 0$ observations, corrected for instrumental degradation over time, and applied exposure corrections that account for Earth's elliptical orbit as well as AIA's AEC. We also have ensured that both AIA and HMI data are spatially colocated, have identical angular resolutions, and that all instruments are chronosynchronous.

We also highlight some of the potential pitfalls of blindly applying machine-learning techniques to solar data, or even more broadly:

1. To maximize its versatility, *SDO* data products are nuanced and assume an expert-level understanding of its instrumental design and limitations. Using them without this knowledge may lead to incorrect results and invalid conclusions.
2. Most of the physical processes that drive solar variability occur at a much slower cadence than that of *SDO*'s instruments (hours and days versus minutes and seconds, respectively), requiring special care with the splitting of training, validation, and test sets. Splits must be performed along temporal blocks and not by random sampling as is done in other settings with uncorrelated data samples. Random sampling in this case will lead to an overly optimistic estimate of validation error, leading to an inability to identify whether a model will generalize properly to future observations or has instead been overfit to its data.
3. Due to the relatively short timescales of solar variability, the simple forecasting models of permanence and climatological averages perform exceptionally well in hourly and daily timescales. Due to this, error estimates of more advanced models are not meaningful in an absolute sense but rather only when compared to these simple baselines.

Finally, we provide a series of baselines that take advantage of this data set to produce EVE time forecasts and HMI-to-AIA

reconstructions. These examples are meant to illustrate some of the applications made possible by combining these data with machine-learning techniques, as well as what heuristic performance measures one should expect to compare their own model implementations with.

As with many fields, heliophysics has entered a data-rich age in which the human intellect alone is incapable of processing the copious amounts of data gathered by NASA's ever-growing spacecraft fleet. Fortunately, the ongoing revolution in machine-learning research will power a new age of data inference and physical insight that maximizes the scientific output of these data-rich missions. It is important, however, for heliophysicists and computer scientists to work together to understand the properties and limitations of both the raw data and the machine-learning techniques. If special care is not taken in understanding such limitations, we may unfortunately see a large amount of incorrect, overly optimistic—or worse—misleading research. Interdisciplinary programs, such as NASA's Frontier Development Laboratory, can provide a vital common ground to facilitate this skill transfer and will be highly critical for the successful and fruitful development of machine-learning techniques in the astrophysical sciences.

This project was conducted during the 2018 NASA Frontier Development Lab (FDL) program, a public/private partnership between NASA and SETI and industry partners including NVIDIA Corporation, Lockheed Martin, and IBM. The authors thank IBM (especially Naeem Altaf) for generously providing computing resources on the IBM Cloud. We gratefully thank our mentors for guidance and useful discussion, as well as the SETI Institute for their hospitality. R.G. acknowledges support from the Moore-Sloan Data Science Environment at New York University and thanks Rob Fergus for useful discussion. The authors acknowledge support from NASA's *SDO*/AIA contract (NNG04EA00C) to LMSAL. AIA is an instrument onboard the *Solar Dynamics Observatory*, a mission for NASA's Living With a Star program. M.C.M.C. acknowledges support from NASA's Heliophysics Grand Challenges Research grant (NNX14AI14G).

Software: PyTorch (Paszke et al. 2017), SunPy (SunPy Community et al. 2015), SolarSoft (Freeland & Handy 1998).












Appendix

The data set is made available through the Stanford Digital Repository, partitioned by year. The data may be obtained via the links provided in Table 3, totaling 6.5TBs.

Table 3
Data Set URLs and Digital Object Identifiers

Year	Stanford Digital Repository url	DOI
2010	https://purl.stanford.edu/vk217bh4910	10.25740/ppax-bf07
2011	https://purl.stanford.edu/jc488jb7715	10.25740/sb4q-wj06
2012	https://purl.stanford.edu/dc156hp0190	10.25740/1vyz-b592
2013	https://purl.stanford.edu/km388vz4371	10.25740/2zme-3q44
2014	https://purl.stanford.edu/sr325xz9271	10.25740/3jhw-x180
2015	https://purl.stanford.edu/qw012qy2533	10.25740/0fbp-re41
2016	https://purl.stanford.edu/vf806tr8954	10.25740/64cr-bc95
2017	https://purl.stanford.edu/kp222tm1554	10.25740/c8bw-ar96
2018	https://purl.stanford.edu/nk828sc2920	10.25740/pknx-5s37

ORCID iDs

Richard Galvez  <https://orcid.org/0000-0002-4780-9566>
 David F. Fouhey  <https://orcid.org/0000-0001-5028-5161>
 Meng Jin  <https://orcid.org/0000-0002-9672-3873>
 Alexandre Szenicer  <https://orcid.org/0000-0002-4829-5739>
 Andrés Muñoz-Jaramillo  <https://orcid.org/0000-0002-4716-0840>
 Mark C. M. Cheung  <https://orcid.org/0000-0003-2110-9753>
 Paul J. Wright  <https://orcid.org/0000-0001-9021-611X>
 Monica G. Bobra  <https://orcid.org/0000-0002-5662-9604>
 Yang Liu  <https://orcid.org/0000-0002-0671-689X>
 James Mason  <https://orcid.org/0000-0002-3783-5509>
 Rajat Thomas  <https://orcid.org/0000-0002-5362-4816>

References

- Bobra, M. G., & Couvidat, S. 2015, *ApJ*, 798, 135
 Boerner, P., Edwards, C., Lemen, J., et al. 2012, *SoPh*, 275, 41
 Boerner, P. F., Testa, P., Warren, H., Weber, M. A., & Schrijver, C. J. 2014, *SoPh*, 289, 2377
 Cao, Z., Simon, T., Wei, S.-E., & Sheikh, Y. 2017, in in Proc. IEEE Conf. Comput. Vision Pattern Recogn. (Piscataway, NJ: IEEE), 7291
 Colak, T., & Qahwaji, R. 2013, *SoPh*, 283, 143
 Freeland, S., & Handy, B. 1998, *SoPh*, 182, 497
 He, K., Zhang, X., Ren, S., & Sun, J. 2016, in in Proc. IEEE Conf. Comput. Vision Pattern Recogn. (Piscataway, NJ: IEEE), 770
 Hoeksema, J. T., Liu, Y., Hayashi, K., et al. 2014, *SoPh*, 289, 3483
 Huang, X., Wang, H., Xu, L., et al. 2018, *ApJ*, 856, 7
 Ioffe, S., & Szegedy, C. 2015, in in Proc. 32nd Conf. Machine Learn. Res., 448
 Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. 2017, in in Proc. IEEE Conf. Comput. Vision Pattern Recogn. (Piscataway, NJ: IEEE), 5967
 Jonas, E., Bobra, M., Shankar, V., Todd Hoeksema, J., & Recht, B. 2018, *SoPh*, 293, 48
 Kucuk, A., Banda, J. M., & Angryk, R. A. 2017, *Scientific Data*, 4, 170096
 Lemen, J. R., Title, A. M., Akin, D. J., et al. 2012, *SoPh*, 275, 17
 Nair, V., & Hinton, G. 2010, in in Proc. 27th Conf. Machine Learn. Res., 807
 Paszke, A., Gross, S., Chintala, S., et al. 2017, in in 31st Conference on Neural Information Processing Systems, Automatic Differentiation in pytorch (Long Beach, CA: NIPS)
 Pesnell, W. D., Thompson, B. J., & Chamberlin, P. C. 2012, *SoPh*, 275, 3
 Scharstein, D., & Szeliski, R. 2002, *Int. J. Comput. Vision*, 47, 7
 Schou, J., Scherrer, P. H., Bush, R. I., et al. 2012, *SoPh*, 275, 229
 Shelhamer, E., Long, J., & Darrell, T. 2016, ITPAM, in press (arXiv:1605.06211)
 SunPy Community, T., Mumford, S. J., Christe, S., et al. 2015, *CS&D*, 8, 014009
 Sutskever, I., Martens, J., Dahl, G., & Hinton, G. 2013, in in Proc. 30th Int. Conf. Machine Learn. Res., 1139
 Wang, X., Fouhey, D. F., & Gupta, A. 2015, in in Proc. IEEE Conf. Comput. Vision Pattern Recogn. (Piscataway, NJ: IEEE)
 Woods, T. N., Eparvier, F. G., Hock, R., et al. 2012, *SoPh*, 275, 115
 Wright, P. J., Cheung, M. C. M., Thomas, R., et al. 2019, in Machine Learning, Statistics, and Data Mining for Heliophysics, ed. M. Bobra, J. Mason et al. <https://helioml.github.io/HelioML>