FlareNet: A Deep Learning Framework for Solar Phenomena Prediction

FDL 2017 Solar Storm Team:
Sean McGregor[1], Dattaraj Dhuri[1,2], Anamaria Berea[1,3], and Andrés Muñoz-Jaramillo[1,4]

[1]*NASA's 2017 Frontier Development Laboratory*
[2]*Tata Institute of Fundamental Research (TIFR), Mumbai, India*
[3]*Center for Complexity in Business, University of Maryland, College Park, MD, USA*
[4]*SouthWest Research Institute, Boulder, CO, USA*

Abstract

Solar activity can interfere with the normal operation of GPS satellites, the power grid, and space operations, but inadequate predictive models mean we have little warning for the arrival of newly disruptive solar activity. Petabytes of data collected from satellite instruments aboard the Solar Dynamics Observatory (SDO) provide a high-cadence, high-resolution, and many-channeled dataset of solar phenomena. Several challenging deep learning problems may be derived from the data, including space weather forecasting (i.e., solar flares, solar energetic particles, and coronal mass ejections). This work introduces a software framework, FlareNet, for experimentation within these problems. FlareNet includes components for the downloading and management of SDO data, visualization, and rapid experimentation. The system architecture is built to enable collaboration between heliophysicists and machine learning researchers on the topics of image regression, image classification, and image segmentation. We specifically highlight the problem of solar flare prediction and offer insights from preliminary experiments.

1 Introduction

The violent release of solar magnetic energy – collectively referred to as "space weather" – is responsible for a variety of phenomena that can disrupt technological assets. In particular, solar flares (sudden brightenings of the solar corona) and coronal mass ejections (CMEs; the violent release of solar plasma) can disrupt long-distance communications, reduce Global Positioning System (GPS) accuracy, degrade satellites, and disrupt the power grid [5].

Predicting space weather is a challenging task because the release of magnetic energy stems from a sudden catastrophic loss of equilibrium in an otherwise meta-stable system (akin to seismological activity or the occurrence of lightning strikes). Current operational space weather relies on hand tailored morphological analyses of the Sun's magnetic field [9], but even ensemble models derived from experts in the field perform close to a persistence baseline [3].

With the launch of the Solar Dynamics Observatory (SDO) [10] in 2010, we have access to a space-based instrument collecting terabytes of full disk solar images on a daily basis. These high-cadence, high-resolution, many-channeled images include maps of the solar magnetic and velocity fields (magnetograms and dopplergrams), as well as images of the solar atmosphere using a variety of wavelengths (see Figure 1 for examples).

The SDO dataset poses unique opportunities and challenges for deep learning. This work introduces "FlareNet" as a deep learning framework for solar physics research to address the research preconditions for modeling space weather. FlareNet includes functionality for data management, neural

Workshop on Deep Learning for Physical Sciences (DLPS 2017), NIPS 2017, Long Beach, CA, USA.

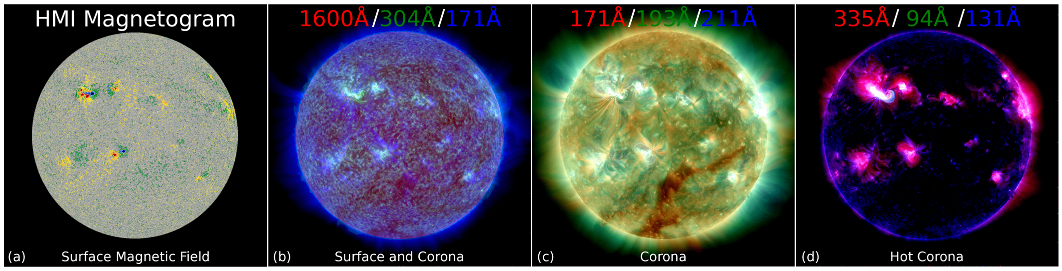

Figure 1: **(a)** The Helioseismic and Magnetic Imager (HMI; [11]) provides 4k X 4k (0.5 arcsec/pixel resolution) full-disk images of surface magnetic field (vector and line-of-sight magnetograms) every 12 minutes and surface velocity (Dopplergrams) every 45 seconds. The yellow and red (green and blue) pixels of the image denote magnetic fields pointing towards (away from) the observer. **(b, c, & d)** The Atmospheric Imaging Assembly (AIA; [7]) captures 8 channels spanning UV and EUV spectrum. These images are also 4k X 4k (0.6 arcsec/pixel resolution), taken at every 12 seconds. Here we composite images captured by AIA for different spectra into the RGB color channels. (a) shows wavelengths observing the surface (red), chromosphere (green), and corona (blue). (b) shows three wavelengths observing the corona. (c) shows three wavelengths observing the hot/active corona. Collectively these images capture the state of visible solar activity.

network specification, training, and visualization of solar phenomena. By formalizing this complete research environment, we can simultaneously leverage the domain knowledge of physical scientists and the neural network architecture experience of computer scientists. During training, a collection of visualization scripts run to help physical scientists interpret the relationships captured by the neural network (see Figure 2). Since the data is fully modeled within FlareNet, deep learning researchers can concentrate on network architectures and avoid the pitfalls of correcting the data for instrument changes and other tradecraft problems.

Our team of computer scientists and heliophysicists developed FlareNet during the 2017 NASA Frontier Development Lab (FDL). We now more fully introduce the physics and the software developed by our team.

2 A Brief Introduction to Solar Physics and FlareNet

The Sun is a hot ball of plasma primarily consisting ionized hydrogen and helium gases. Dark spots called sunspots, which are relatively cooler areas, appear on the surface with their number and surface area varying through 11 year solar cycles [6]. Sunspots are surrounded by regions of concentrated magnetic field called active regions. Magnetic field activity produced inside the sun follows plasma motion “flux tubes” to the solar surface. Magnetic flux tubes are stretched and twisted by plasma motion and reach into the solar atmosphere to form giant loop structures over active regions [14]. These magnetic loops store energy. As magnetic fields rise to the surface and into the solar atmosphere, energy builds up and occasionally releases in eruptions such as solar flares [12].

SDO monitors solar activity and eruptive events with several space-based instruments (see Figure 1). This high resolution SDO dataset poses unique challenges for deep learning. Each pixel exhibits very high dynamic ranges with flux that tends to confuse gradient updates and encourage overfitting. FlareNet addresses these, and other issues to make the problem more amenable to deep learning.

We built FlareNet with components for downloading and transforming SDO data, specifying network architectures [4], and running experiments. During training, a collection of visualization scripts run to help physical scientists interpret the relationships being captured by the neural network. Physical scientists can enhance the understanding process by contributing additional visualization scripts (see Figure 2).

We also incorporated several useful tools for modifying FlareNet inputs. First, in traditional video processing techniques it is necessary to incrementally construct a model of the state by sequentially processing multiple time steps, but this is not necessarily required for solar images. For high-cadence, scaled and centered data, temporally adjacent images capture the same spatial locations and we can treat the time steps as additional image channels. FlareNet supports this “temporal compositing”

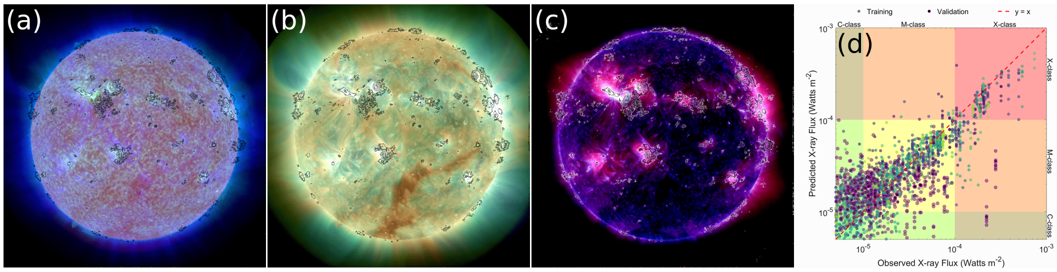

Figure 2: Saliency map overlaid over the AIA composites showing the surface+chromosphere+corona (a), corona (b), and hot corona (c). (d) Observed vs. forecasted flare X-ray flux for training (2010-2015) and validation (2016) sets.

through the dataset API. Second, the data API supports defining vectors of side channel information [15] to be appended to the fully connected layers of the neural network. These side channels allow the convolution layers to avoid learning concepts like the 11 year solar cycle by directly providing the information. Finally, we define optional pre-processing layers for log transforming the input images on the GPU. These transformations help address the challenges introduced by high dynamic range SDO images.

3 Space Weather Task Definitions

The problems of solar flare, irradiance, and CME forecasting can all be formulated as image classification, image regression, pixel regression (predicting real valued output at individual pixels), and pixel segmentation (discretized pixel regression) problems within FlareNet. We can also formulate the problem of solar particle emissions as classification and regression problems, but particle measurement does not ascribe solar outputs to individual pixels, thereby preventing pixel regression.

Each of these tasks share the same set of independent variables (the SDO images). We support the classification and regression tasks by changing between files mapping time steps to the dependent variable.

Predicting each of these phenomena within FlareNet requires specifying two task metaparameters. These include the *lag* until the prediction will be made, and the *time frame* within which we predict the phenomena. Setting the time lag to larger values is similar to issuing an extended forecast. Setting the time frame to larger windows tends to smooth the noise of solar phenomena and increase the probability of capturing events.

We highlight the problem of solar flare prediction as an illustrative case for the image regression, classification, and segmentation problems posed by the sun. Solar flares travel from the sun to earth at the speed of light, which means that we have no warning before their electromagnetic radiation interacts with our atmosphere. Predicting solar flares is challenging because their triggering mechanism exhibits similar behavior to snow avalanches or earthquakes [13] – a property known as "self-criticality" [2, 8]. Free-energy available for this process arises from the organization or the solar magnetic field, and thus the current methodology for flare forecasting (based on a morphological classification of magnetic regions [9]) can be understood as an quantification of the storage of free magnetic energy.

Although the sun produces low magnitude flares with high frequency, only the strongest flares produce observable adverse technological impact. Flare magnitude is typically tied to the amplitude of their X-ray radiative output as measured by the GOES X-Ray satellites [1] (see Fig. 3-a). In this work we focus on flares of class C, M, and X as specified by NOAA's flare catalog.[1]

Considering that the SDO era (2010-2017) has only around 8,000 flares of class equal or greater than C-class, the 12 second cadence of SDO images leads to a class imbalance between flaring and non-flaring images. Thus, naive training makes the network trend towards always predicting the

[1]This catalog contains the X-ray flux, start, peak and end times for most flares observed between 1975 and 2017. https://www.ngdc.noaa.gov/stp/solar/solarflares.html

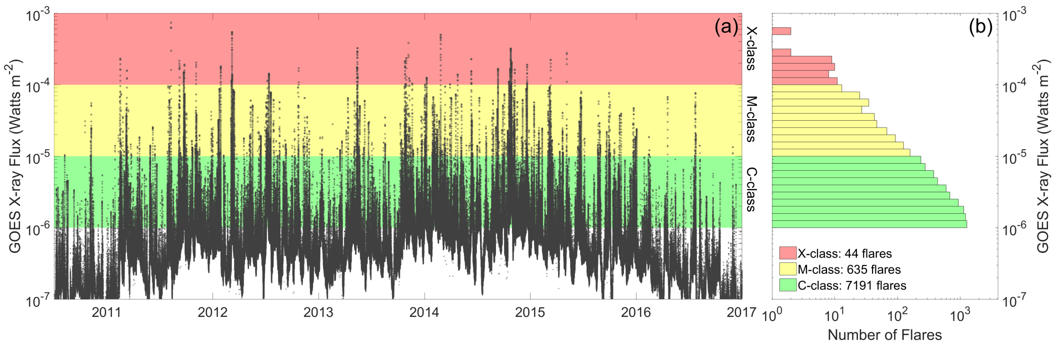

Figure 3: (a) X-Ray flux measured by the GOES satellite during the SDO era at a cadence of 2 minutes. Flare classification is done based on the maximum X-ray flux emitted during the flaring event. The strongest category (X-class; red) contains flares ten times stronger than the class below it (M-class; yellow) and 100 times stronger than the next one (C-class; green). (b) Histogram of flare occurrence during the SDO era. A clear power law can be seen in the flare distribution showing that M-class (C-class) flares are roughly 10 (100) times more numerous than X-class flares.

persistence baseline. Because of this, we introduced several oversampling strategies to FlareNet, including limiting training to pre-flaring images. This strategy steps away from the original task of forecasting when a flare would happen and how strong it would be, to focus on how strong a flare could be given the current state of the Sun. Within the avalanche metaphor, this approach is akin to measuring the amount of snow on the mountain.

4 Discussion

We developed FlareNet during a six week intensive collaboration. Our interdisciplinary team of computer scientists and physicists lacked sufficient training time to iterate on network architectures, but our software and problem definition offer important insights into next steps for addressing space weather problems. FlareNet supports several additional solar modeling problems by changing the dependent variables within FlareNet to the CME catalog, irradiance measurements, or particle emissions. It is our hope that by putting in the architectural effort required to develop FlareNet, we might inspire additional cross-disciplinary collaboration in the physical sciences.

Acknowledgments

We are very grateful to IBM and IBM's Troy Hernandez and Naeem Altaf for their incredible support in terms of hardware and tech. Without their support this work would not have been possible. We are very grateful to Kx, and to NASA's Heliophysics division for providing financial support that enabled this interdisciplinary collaboration and to the SETI Institute for providing the logistic framework that nurtured and enabled our research.

We are grateful to Mark Cheung, Monica Bobra, Ryan McGranaghan, and Graham Mackintosh, co-mentors at NASA's Frontier Development Laboratory, for valuable suggestions and discussions.

The X-ray Flare dataset was prepared by the NOAA National Geophysical Data Center (NGDC). SDO data were prepared and made available for download by Stanford's Joint Science Operations Center (JSOC).

References

[1] Markus J Aschwanden. Irradiance observations of the solar soft x-ray flux from goes. In *The Sun as a Variable Star: Solar and Stellar Irradiance Variations*, pages 53–59. Springer, 1994.

[2] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality - An explanation of 1/f noise. *Physical Review Letters*, 59:381–384, July 1987.

[3] G. Barnes, K. D. Leka, C. J. Schrijver, T. Colak, R. Qahwaji, O. W. Ashamari, Y. Yuan, J. Zhang, R. T. J. McAteer, D. S. Bloomfield, P. A. Higgins, P. T. Gallagher, D. A. Falconer, M. K. Georgoulis, M. S. Wheatland, C. Balch, T. Dunn, and E. L. Wagner. A comparison of flare forecasting methods. i. results from the "all-clear" workshop. *The Astrophysical Journal*, 829(2):89, 2016.

[4] François Chollet and Others. Keras. https://github.com/fchollet/keras, 2015.

[5] J. P. Eastwood, E. Biffis, M. A. Hapgood, L. Green, M. M. Bisi, R. D. Bentley, R. Wicks, L.-A. McKinnell, M. Gibbs, and C. Burnett. The economic impact of space weather: Where do we stand? *Risk Analysis*, 37(2):206–218, 2017.

[6] David H. Hathaway. The solar cycle. *Living Reviews in Solar Physics*, 7(1):1, Mar 2010.

[7] 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 Auker, Paul Jerram, Peter Pool, Regina Soufli, 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):17–40, Jan 2012.

[8] E. T. Lu and R. J. Hamilton. Avalanches and the distribution of solar flares. *ApJL*, 380:L89–L92, October 1991.

[9] P. S. McIntosh. The classification of sunspot groups. *Sol. Phys.*, 125:251–267, February 1990.

[10] W. D. Pesnell, B. J. Thompson, and P. C. Chamberlin. The Solar Dynamics Observatory (SDO). *Sol. Phys.*, 275:3–15, January 2012.

[11] J. Schou, P. H. Scherrer, R. I. Bush, R. Wachter, S. Couvidat, M. C. Rabello-Soares, R. S. Bogart, J. T. Hoeksema, Y. Liu, T. L. Duvall, D. J. Akin, B. A. Allard, J. W. Miles, R. Rairden, R. A. Shine, T. D. Tarbell, A. M. Title, C. J. Wolfson, D. F. Elmore, A. A. Norton, and S. Tomczyk. Design and ground calibration of the helioseismic and magnetic imager (hmi) instrument on the solar dynamics observatory (sdo). *Solar Physics*, 275(1):229–259, Jan 2012.

[12] Kazunari Shibata and Tetsuya Magara. Solar flares: Magnetohydrodynamic processes. *Living Reviews in Solar Physics*, 8(1):6, Dec 2011.

[13] A Sornette and D Sornette. Self-organized criticality and earthquakes. *EPL (Europhysics Letters)*, 9(3):197, 1989.

[14] Robert F. Stein. Solar surface magneto-convection. *Living Reviews in Solar Physics*, 9(1):4, Jul 2012.

[15] Yilun Zhou and Kris Hauser. Incorporating Side-Channel Information into Convolutional Neural Networks for Robotic Tasks. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on.* IEEE, 2017.