The *SDO*/EVE Solar Irradiance Coronal Dimming Index Catalog. I. Methods and Algorithms

James Paul Mason[1,2] , Raphael Attie[1] , Charles N. Arge[1], Barbara Thompson[1] , and Thomas N. Woods[2]

[1] NASA Goddard Space Flight Center, 8800 Greenbelt Rd., Greenbelt, MD 20771, USA; james.p.mason@nasa.gov

[2] Laboratory For Atmospheric and Space Physics, University of Colorado at Boulder, 3665 Discovery Dr., Boulder, CO 80303, USA

james.mason@lasp.colorado.edu

Abstract

When a coronal mass ejection departs, it leaves behind a temporary void. That void is known as coronal dimming, and it contains information about the mass ejection that caused it. Other physical processes can cause parts of the corona to have transient dimmings, but mass ejections are particularly interesting because of their influence in space weather. Prior work has established that dimmings are detectable even in disk-integrated irradiance observations, i.e., Sun-as-a-star measurements. The present work evaluates four years of continuous *Solar Dynamics Observatory* Extreme Ultraviolet Experiment (EVE) observations to greatly expand the number of dimmings we may detect and characterize, and collects that information into James's EVE Dimming Index catalog. This paper details the algorithms used to produce the catalog, provides statistics on it, and compares it with prior work. The catalog contains 5051 potential events (rows), which correspond to all robustly detected solar eruptive events in this time period as defined by >C1 flares. Each row has a corresponding 27,349 elements of metadata and parameterizations (columns). In total, this catalog is the result of analyzing 7.6 million solar ultraviolet light curves.

1. Introduction

Coronal dimming is a phenomenon that has been observed for decades, starting at least as far back as Skylab X-ray observations (Rust & Hildner 1976). Originally referred to in the literature as "transient coronal holes," subsequent observations were made both in soft X-ray (e.g., Hudson et al. 1996) and extreme ultraviolet (EUV) image data (e.g., Thompson et al. 2000; Reinard & Biesecker 2008). In recent years, work has been done to establish a connection between coronal mass ejection (CME) kinematics and dimming observed in these EUV images (Aschwanden 2009; Dissauer et al. 2018a, 2018b). A surprising result from the *Solar Dynamics Observatory* (*SDO*; Pesnell et al. 2012) EUV Variability Experiment (EVE; Woods et al. 2010) is that dimming is so pronounced that it impacts the total solar energy output in some EUV spectral emission lines (Woods et al. 2011), i.e., dimming is measurable in irradiance, which by definition has no spatial resolution. Mason et al. (2014) studied these Sun-as-a-star measurements in conjunction with EUV images of dimming and white-light coronagraph images of an associated CME. They found that while the dimming could be spatially isolated in EUV images to produce a light curve representing only the dimming region, it was also possible to produce very similar light curves using only the spatially integrated *SDO*/EVE data. Mason et al. (2016) then applied this dimming-isolation method to 37 events to establish a statistical relationship between irradiance coronal dimming depth and CME mass and between dimming slope and CME speed. Dissauer et al. (2018a) performed a similar study of 62 events using the *SDO* Atmospheric Imaging Assembly (AIA; Lemen et al. 2012) EUV

imager to identify dimming and found similar dimming-CME empirical relationships as those in Mason et al. (2016). This paper represents the first step in the natural extension of the case study in Mason et al. (2014) and semi-statistical study in Mason et al. (2016): the automation of dimming detection and characterization in the entire history of *SDO*/EVE Multiple EUV Grating Spectrographs A (MEGS-A) data. That automation required that key times be input as triggers to search for dimming. The most long-standing and robust identifier of solar eruptive events is the *Geostationary Operational Environmental Satellite* system (*GOES*) X-ray Sensor (XRS) flare event list. Flares, CMEs, SEPs, and dimmings are all separate processes, but all are manifestations of the rapid magnetic energy release that produces solar eruptive events. Some CMEs may occur with no measurable dimming, and some dimmings may have no observed and/or measurable CMEs. Some CMEs occur without an associated flare, though these tend to be small and slow (Yashiro et al. 2005; Ma et al. 2010), which would presumably also result in dimmings that are correspondingly difficult to measure. Thus, we chose to generate this first iteration of the catalog using flares as the key times triggering searches for dimming. Given the catalog that results from this work, it will be possible to compare to large CME catalogs—e.g., CDAW (Gopalswamy et al. 2009) and CACTus (Berghmans et al. 2002)—in order to derive statistically robust relationships and to identify and study outliers.

This work will be broken up into at least two papers. The present one describes the methods and algorithms used to produce James's EVE Dimming Index (JEDI) catalog (Mason 2019). Subsequent papers will describe comparisons with other dimming catalogs—e.g., Solar DEMON (Kraaikamp & Verbeeck 2015) and CoDiT (Krista & Reinard 2017)—and the previously mentioned CME catalogs. Specifically, this paper describes the input and output of the JEDI pipeline in Section 2 and the methods and algorithms applied to the input to produce the output

in Section 3. Section 3.1.1 describes a sensitivity study done on a crucial component of the overall algorithm: pre-flare irradiance determination, i.e., establishing the baseline from which dimming is measured. Section 4 lists the throughput percentages of each component of the algorithm and then provides some basic statistics of the catalog itself. Section 5 pulls out some specific examples from the catalog and compares those results to prior work. Finally, a summary and discussion of future work are provided in Section 6.

2. Input/Output

The main inputs are the *SDO*/EVE data and the *GOES*/XRS flare event list. The main output is the JEDI catalog itself, a large table in the Hierarchical Data Format (HDF) 5 file.

2.1. Inputs

For this work, we use the *SDO*/EVE level 2 extracted emission lines product. There are 39 emission lines extracted from the spectra. Table 1 lists these extracted lines, and Figure 1 shows all of the data used to produce this catalog. Note that Figure 1 uses a linear color intensity scale, so very large dynamics (e.g., flares) are saturated. Additionally, the natural data cadence of 10 s cannot easily be represented in a single figure spanning 4 yr, and as a result, data gaps can be obscured. In particular, the data that are measured by the MEGS-B component of EVE (35–105 nm) were reduced from continuous coverage to a few hours of exposure per day after 2010 May 19. The data began on 2010 May 1 but were truncated to 2014 May 26 because EVE/MEGS-A had a CCD camera power anomaly that could not be resolved. MEGS-B observations continue to the present date but at the reduced cadence not trivially amenable to coronal dimming analysis. The upshot is that JEDI's primary input is a static data set (barring version updates to the data set), which makes comparisons between future versions of JEDI simpler. Table 1 and Figure 1 show that the *SDO*/EVE extracted lines span three orders of magnitude in temperature, with Fe ions sensitive to much of that range. Some ions appear at multiple wavelengths, but these have the same peak formation temperature. We expect similar temperature light curves to have similar trends, though photons at very different wavelengths will be subject to different absorption cross sections in line-of-sight plasma.

The other input to JEDI is the *GOES* flare event list, which was obtained using the IDL SolarSoft routine, *rd_gev*. These events are used as key times to trigger a search for dimming in the *SDO*/EVE data and also to identify the end time for pre-flare irradiance determination windows (described in Sections 3.1 and 3.1.1). For the initial version of JEDI, we have limited the flare events to those above a C1 classification because "big flare syndrome" (Kahler 1982) suggests that the smaller flares are also likely to be subtler in all indicators of energetic release, including dimming. Each event includes the flare class, start and peak times, measured flux, solar latitude and longitude, and an identifier for the associated active region responsible for the flare. We convert the latitude and longitude to the position angle commonly used for CMEs (Equation (1), see Equations (19) and (1) in Thompson 2006[3]). This is done to make future comparisons with CME catalogs easier:

$$\psi = \arg\left(\sin(\text{lat}),\ -\cos(\text{lat})\sin(\text{lon})\right). \qquad (1)$$

---

[3] Edge cases must be handled separately. See *lat_lon_to_position_angle.py* in Mason (2018).

Table 1
SDO/EVE Extracted Emission Lines

Ion	λ (nm)	$\log_{10}T$ (K)	Blends
He I	53.70	3.84	⋯
H I	94.97	3.84	⋯
H I	97.25	3.84	⋯
H I	102.57	3.84	⋯
He I	58.43	4.16	⋯
O II	71.85	4.48	O II
O II	83.55	4.52	⋯
He II	30.38	4.70	⋯
He II	25.63	4.75	⋯
C III	97.70	4.84	⋯
O III	52.58	4.92	O II
O III	59.96	4.92	⋯
O IV	55.44	5.19	O IV
O IV	79.02	5.19	O III, O IV
O V	62.97	5.37	⋯
O VI	103.19	5.47	⋯
Fe VIII	13.12	5.57	Fe VIII
Ne VII	46.52	5.71	⋯
Fe IX	17.11	5.81	⋯
Ne VIII	77.04	5.81	⋯
Fe X	17.72	5.99	Fe VII
Mg IX	36.81	5.99	Fe VIII–Fe XIV
Mg X	62.49	6.05	⋯
Fe XI	18.04	6.07	Fe X, Fe VII
Mg X	60.98	6.10	⋯
Fe XII	19.51	6.13	⋯
Fe XIII	20.20	6.19	Fe XI, Fe XII
Fe XV	21.13	6.27	Fe XII
Si XII	52.10	6.28	⋯
Si XII	49.94	6.29	⋯
Fe XV	28.42	6.30	⋯
Fe XVI	33.54	6.43	⋯
Fe XVI	36.08	6.43	⋯
S XIV	44.57	6.44	⋯
Fe XVIII	9.39	6.81	Fe IX, Fe X
Fe XIX	59.22	6.89	⋯
Fe XX	56.79	6.96	⋯
Fe XX	72.16	6.96	⋯
Fe XX	13.29	6.97	Fe XXIII

Note. Line blends occur where multiple emission lines are very close in wavelength.

2.2. Output

The JEDI catalog is a large HDF5 file. In version 2, there are 5051 rows by 27,349 columns. Each row represents a different potential dimming event and corresponds to a particular *GOES* flare event key time. The columns include several element types. The primary elements are the parameterizations of dimming: depth, slope, and duration (see Sections 3.4–3.6) for each of the 39 EVE extracted emission lines and every permutation of those lines combined by subtraction via the method described in Section 3.2. This is the main reason that there are so many columns: the permutation $n = 39$, $r = 2$ with no repetition is 1482. Adding that to the unaltered 39 emission lines, we end up with 1521 light curves for every potential event. In each of those light curves, we attempt to detect and characterize dimming. The remaining columns in JEDI are essentially metadata. These include timestamps associated with the dimming parameterizations, the relevant *GOES* event information, a flag indicating whether another flare interrupted the dimming characterization, and fitting
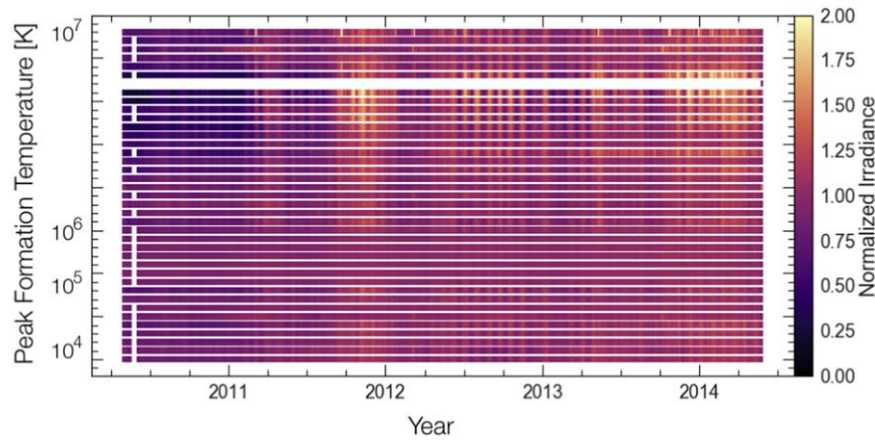
Figure 1. All *SDO*/EVE data used to produce the JEDI catalog. Each horizontal strip corresponds to an emission line as listed in Table 1. The vertical white stripe in many lines in 2010 is the result of a thermal bakeout of MEGS-B CCD, resulting in a gap in solar observations. The horizontal white stripe is from the Fe XVI 36.08 line, which is extracted but always set to the invalid flag. The vertical bright stripes are due to solar rotation variability, such as the coming and going of active regions.

scores. Finally, the processing algorithm also generates numerous plots at each step, including a summary of the final result, an example of which is shown in Figure 2.

It is important to note that the JEDI catalog in its present state does not actually determine whether or not there is dimming for any particular event. Instead, it provides numerous quantitative measurements of light curves. Whether or not those measurements are indicative of dimming depends on the thresholds applied to those measurements, e.g., a user might decide that any event with a 17.1 nm dimming depth >1% is a detected dimming. Future papers will identify good metrics for flagging dimming in JEDI that are internally consistent and compatible with other catalogs.

The catalog is not intended to be perused by eye to search for patterns. It is much too large. Instead, the structure of the JEDI catalog was directly influenced by common formats used in machine learning, which is specifically designed to handle large tables such as this. The file can be easily loaded as a (events × features) Pandas DataFrame, which is the preferred format of the scikit-learn machine-learning package. Additionally, despite the input being small compared to most statistical image-based analyses, the processing time to produce the catalog is much greater than the time to discard unwanted columns by slicing or machine-learning techniques that discard redundant or low-value information, such as principal component analysis.

3. Algorithm

The algorithm to convert *SDO*/EVE emission line data and *GOES* flare events into the JEDI catalog consists of several steps, each of which are described in a separate subsection below. They are pre-flare irradiance determination to establish a baseline from which the depth of dimming can be measured (Section 3.1), light-curve flare removal (Section 3.2), light-curve fitting to minimize the influence of noise (Section 3.3), and determination of dimming depth (Section 3.4), slope (Section 3.5), and duration (Section 3.6).

3.1. Baseline Determination

In order to measure dimming depth in light curves (Section 3.4), a background level must be established from which we can measure that relative decrease. Determining that
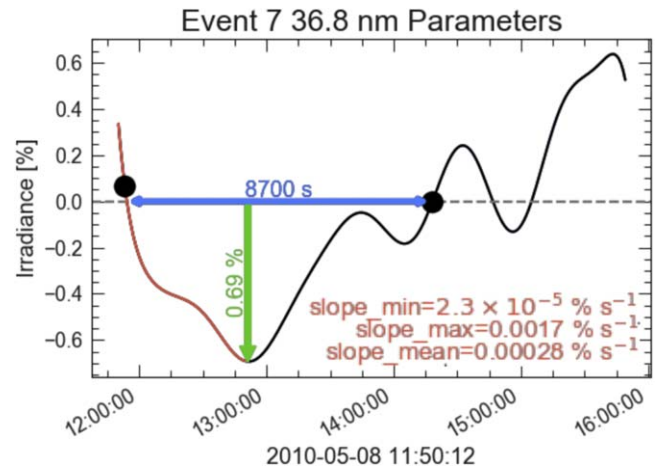

Figure 2. Example summary plot from the JEDI catalog. Green indicates the dimming depth, red the period for slope calculations, and blue the duration.

background baseline level is the objective of this first step in the algorithm.

The baseline should not be increased due to the presence of flares. As described by Ryan et al. (2012), a flare light curve consists of a background flux plus a quiescent flux (a.k.a., the pre-flare flux) and a flare flux that they define as the flux above background. The baseline established here corresponds to their pre-flare flux definition. In other words, if the flare did not occur, what would the flux be? Because every flare is different and their impact on different wavelengths can vary, the simplest way to answer this question is to look at pre-flare times and assume that the flux level would have continued for the subsequent few hours.

The first step in this algorithm is to define a pre-flare time window in which to attempt calculating an irradiance. There are no obvious values for the length of this window. The clearest constraints are that it has to be far enough back in time to include at least one point of data from the instrument (*SDO*/EVE in this case) but not so far back in time that it encompasses prior flares. With only a single point, the baseline could be unduly influenced by a random fluctuation. Spanning too far back in time may be influenced by active region evolution. Clearly, the definition of this window length is not

trivial; thus, it is the subject of a sensitivity study detailed in Section 3.1.1. The result of that study is that the optimal window duration is 300 minutes.

The second step is to divide the pre-flare time window into three equal sub-windows. Within each, the median and standard deviation are calculated. In the ideal case, the light curve would be perfectly flat across the whole window, resulting in identical median and standard deviation (σ) values in each sub-window. Instead, if the light curve shows some trending, we can only assume that the trend would have continued if not for the flare and potential dimming. Indeed, that trend may still be present despite these transient events, but the effects will all be superimposed. If the trend was due to the same active region responsible for the flare and dimming, then we could reasonably expect that the trend should be altered in some manner that is nontrivial to predict. Thus, if we detect such a substantial trend, we simply throw out the whole event. This is a common occurrence: the throughput of this part of the algorithm is only 30.49%. See Section 3.1.1 for the method to determine that value and Section 4 for a listing of the throughput of each subsequent step in the algorithm.

What comprises a "substantial" pre-flare trend? This is where we apply conditions on the median and σ values in the sub-windows. Figure 3 panels (a) and (b) show some examples of basic trends that we want to throw out, and panel (c) shows a limiting case that we would want to accept. In all cases, the programmatic condition applied to determine whether a light curve will be accepted is

$$\max(\mu_n - \mu_m) < T_\mu \times \text{mean}(\sigma_n) \qquad (2)$$

$$\sigma_{n,m} < T_\sigma \qquad (3)$$

where $\max(\mu_n - \mu_m)$ computes the difference between each sub-window's median and selects the largest one, T_μ is a specified threshold for that max median difference that is then scaled by the average "noisiness" of the sub-windows represented by $\text{mean}(\sigma_n)$, and $\sigma_{n,m} < T_\sigma$ means at least two of the three windows must have standard deviations smaller than a specified threshold.

In the examples shown in Figure 3, it can be seen that the light curve in panel (a) will be rejected because both conditions will fail. This concurs with what would be expected when judging this light curve by eye, because it clearly shows a downward trend that may or may not continue during the flare and any potential dimming. The light curve in panel (b) fails the σ condition, and even though it passes the median condition, this light curve would be rejected. Again, this is in line with a subjective judgement because the light curve is clearly oscillatory and that trend may or may not continue. The light curve in panel (c) is an example of a nonideal light curve that would be allowed to pass. One window, w2 in this case, has a very large σ. The median values are all very similar to each other, so that condition passes. The σ condition only requires that two of the three sub-windows have small standard deviations, so even though $\sigma_{w2} > T_\sigma$, a baseline will be computed. The values of T_μ and T_σ were determined by iterative tuning. The results for hundreds of light curves were used and subjectively judged until the algorithmic conditions conformed to subjective expectations; namely that a light curve that would be subjectively hard to identify a clear pre-flare baseline from would also be rejected by the algorithm. Those values are $T_\mu = 1.5\%$ and $T_\sigma = 1\%$.
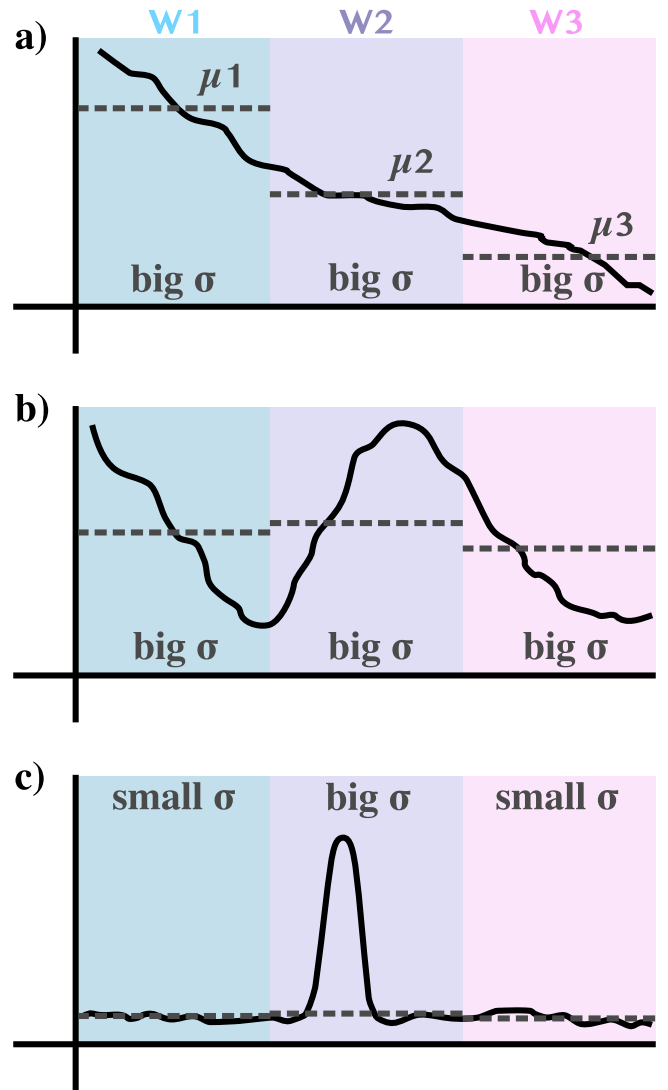

Figure 3. Cartoon examples of pre-flare light curves and sub-windows. μ in this context refers to median rather than mean.

For light curves that have passed these conditions, the computed baseline irradiance is simply the mean of the three medians already calculated. Figure 4 shows two real examples: one that passes the conditions and one that fails. The noisiness of the real light curves is apparent here. In the light curve that passes the conditions, the underlying trend barely extends beyond the noise, while the trend clearly dominates the noise in the light curve that failed the conditions. Indeed, while σ is fairly small in each sub-window, the difference of the medians is substantial.

3.1.1. Sensitivity Study of Pre-flare Irradiance Determination Method

The baseline threshold conditions, T_μ and T_σ, were manually tuned by looking at their impact on hundreds of light curves until the desired result was obtained, i.e., light curves with a clearly identifiable stable background were captured by the algorithm. The remaining free parameter in this determination is the total window size in which to search for a baseline. In terms of limits, a window that encapsulates only a single data point is likely to be unduly influenced by noise. The time window also cannot be infinitely wide because the corona changes significantly over periods of more than a few hours or
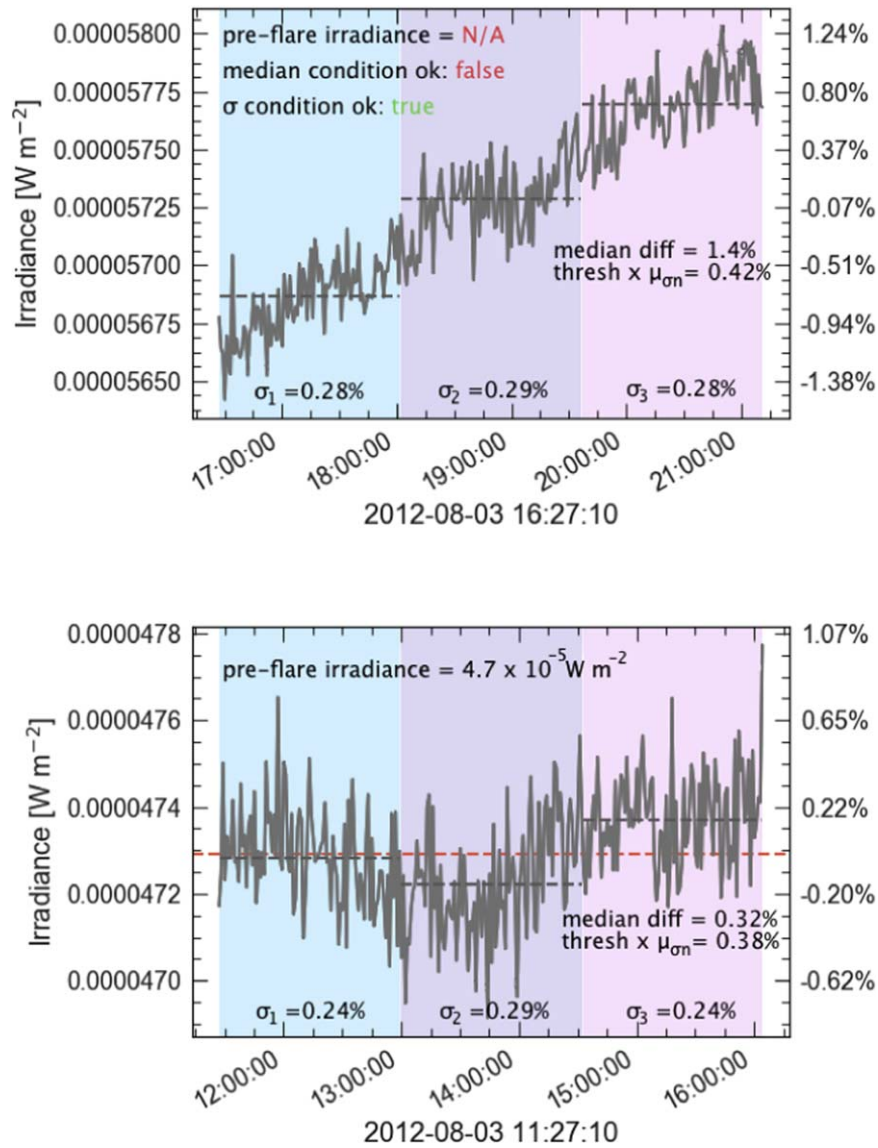
Figure 4. Real examples of pre-flare light curves that fail (top) and pass (bottom) the sub-window conditions. Percentages on the right axis are computed from the median irradiance in the window.

days. Furthermore, if the events are closely spaced together, long windows can result in overlap. In these situations, we propagate forward the first "original" baseline irradiance. With an exceptionally long window, this overlap occurs more often, and it becomes more likely that the baseline level for later events is "stale." This motivated a sensitivity study to determine what window duration resulted in the highest throughput of the baseline determination algorithm.

The throughput is computed as the number of successful pre-flare irradiance determinations divided by the total number of attempts. The latter value is not strictly a fixed value because of the overlap forward propagation of "original" baseline irradiance as discussed above. Table 2 shows the results of this analysis.

The best performing window duration is 300 minutes with a throughput of 30.49%. Even in this best case, 70% of the light curves failed to pass the baseline threshold conditions. This is consistent with "by eye" baseline determinations; it is often hard to identify a baseline with confidence, and spot checks of

Table 2
Baseline Determination Throughput as a Function of Time Window Size

Window (minutes)	# Successful	# Failed	Total	Throughput (%)
60	40536	119832	160368	25.28
120	26764	88598	115362	23.20
180	25519	59657	85176	29.96
240	20194	47354	67548	29.90
300	16804	38303	55107	30.49
360	13710	31686	45396	30.20
420	11204	26626	37830	29.62
480	9708	23910	33618	28.88
720	5808	15447	21255	27.33
1440	1913	6784	8697	22.00

the results (e.g., Figure 4) showed that the algorithm was accepting/rejecting light curves appropriately. Additionally, the algorithm does not appear to be especially sensitive to the window duration, with throughput spanning only 7 percentage

points across windows ranging from a single hour to an entire day. Nevertheless, we chose 300 minutes as the default value to maximize throughput. A minor but pragmatic concern is processing time. A smaller window results in more events to process. The time to process scales linearly with the number to process, so we have a small preference for longer windows. Balancing this preference against the concern with stale forward propagation, we would naturally gravitate toward a moderate window duration. Fortunately, this also aligned with the maximum throughput.

3.2. Flare Removal

Any time there is a dimming, it is likely that there is also a solar flare around the same time, especially as the dimmings become more intense and thus more easily measurable. This temporal coupling can be segregated spatially with EUV imagers: the bright flare pixels can be isolated from surrounding and nearby dimming. Fortunately, the strong spatial distinction is generally also a strong temperature distinction. The volumes that get very bright also tend to be where plasma is getting very hot. Even though irradiance instruments have no spatial resolution, they can distinguish between temperatures if they can spectrally resolve temperature-sensitive emission lines. This is the case with SDO/EVE. Its 39 extracted emission lines span three orders of magnitude in temperature (Table 1 and Figure 1). The light curves of these various emission lines trace the temperature evolution of the plasma, which is most dramatic during solar flares. Emission lines from ions with high peak formation temperatures (e.g., Fe XX at nearly 10 MK) are very weak most of the time but experience a major enhancement from the localized plasma in a solar flare. The flare contributions for the hotter coronal emissions are usually a short-duration impulsive phase and a longer-duration gradual phase afterwards (Woods et al. 2011). The resultant light curves show a flat baseline, a rapid peak, and a quick return to the flat baseline. At temperatures closer to those of the ambient corona, $\sim$1 MK, a larger volume of plasma participates in the emission. Light curves from emission lines near these ambient temperatures (e.g., Fe IX at 0.6 MK) display a corresponding sensitivity to more processes in the corona. These light curves do still show an enhancement from solar flares, though only at later times once the hot, localized plasma has had time to cool. These light curves can also present dimming. The net result is an irradiance light curve that superimposes all of these processes. The dimming may dominate at one point in time, then be overtaken by the solar flare enhancement for a time, return to a dimming level, and eventually return to a baseline level.

With some light curves sensitive to both dimming and flares and other light curves sensitive only to flares, it is possible to combine them to remove the effect of the flare in the former, thus isolating the dimming. This procedure was first outlined in Mason et al. (2016) and has been generalized here. The algorithm is simple in concept: (1) find the solar flare peak in both light curves, (2) time shift the second light curve so that its peak is at the same time as the first, (3) scale the second light curve so that its peak value matches that of the first, and (4) subtract the second light curve from the first. Note that the algorithm will reject light curves if it cannot find peaks or if the required time shift is greater than 30 minutes—corresponding to the longest time it typically takes for plasma to cool across the SDO/EVE emission lines. Figure 5 shows an example of
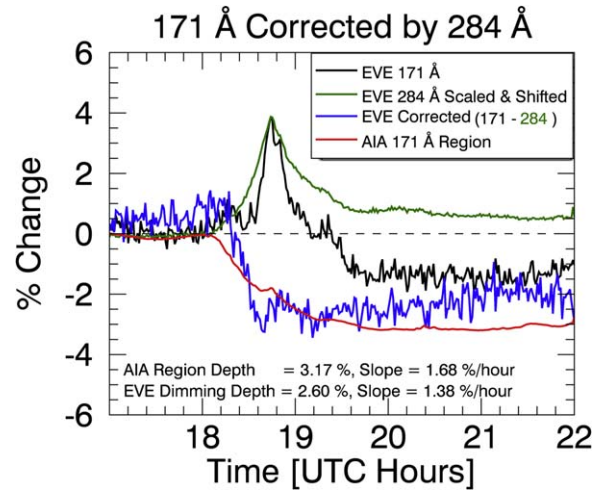

Figure 5. Irradiance flare removal algorithm compared to spatial isolation of dimming, adapted from Mason et al. (2016).

this process and compares it to the corresponding spatially isolated dimming light curve from SDO/AIA. Mason et al. (2016) did this for a handful of emission line combinations and found that the dimming light curve is often best in the cooler, ambient corona lines (e.g., Fe IX 17.1 nm) and that the reference flare light curve is often best in the hotter coronal lines (e.g., Fe XV 28.4 nm) as is shown in Figure 5. The present work does this for every possible permutation, which results in 1482[4] new combined light curves. While we do not expect all of these to be valuable, we avoid introducing any a priori bias in the "right" lines to combine. Subsequent steps of the overall algorithm use these light curves in addition to the original 39, bringing the total number of light curves to process for each event to 1531.

3.3. Light-curve Fit

In order to reduce the influence of minor variations in the light curves on the subsequent dimming measurements, each of the 1531 light curves for each potential event are fit with a smooth curve. The first step of the fitting algorithm is to define the time window in which dimming will be characterized. We start the window one minute prior to the time of the flare peak and end it 24 hr after the flare peak time. This time window was chosen in accordance with the findings from Zarro et al. (1999) and Thompson et al. (2000) that dimming typically lasts several hours and very rarely longer than 1 day. If another flare occurs during this time, the window is truncated to the time of that subsequent flare. If that makes the window less than 2 hr, we reject the light curve. It is likely that some of these events could be assessed for dimming manually, but automated dimming characterization in these cases would be difficult.

With the window established, we can fit a curve to the data. Mason et al. (2016) did this using polynomials and showed that in addition to creating a smoother curve, the uncertainties were reduced. In the present work, we use a machine-learning technique to find the best-fit curve to achieve the same basic result. For each light curve, a series of support vector machine regressions (SVRs; Drucker et al. 1997) from *scikit-learn* (Pedregosa et al. 2011) are applied and evaluated. The primary tunable parameter in SVR is γ, analogous to order in
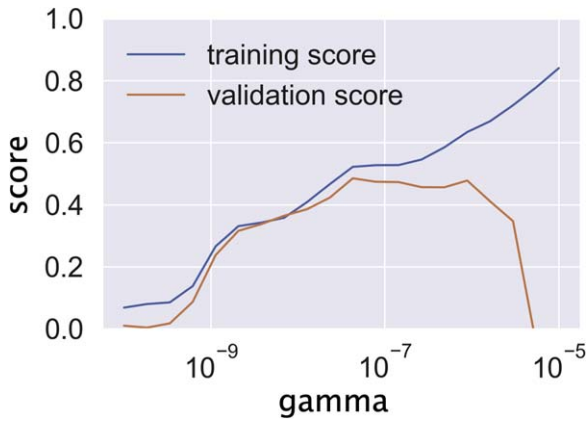
---

4 $\frac{39!}{(39-2)!} = 1492$

Figure 6. Example validation curve for finding the best fit to the light curve.
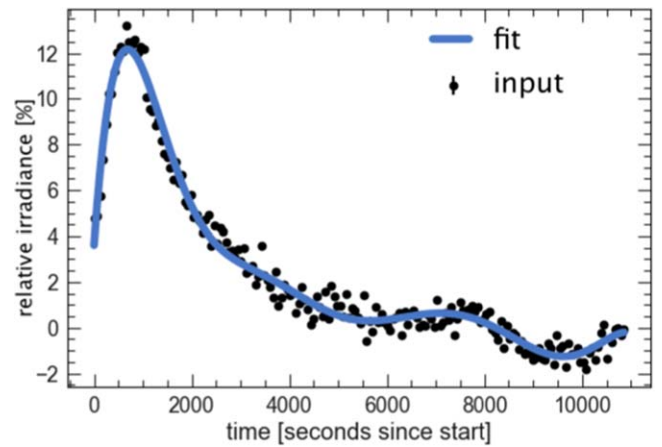

Figure 7. Example light-curve fit for 9.4 nm line with $\gamma = 5 \times 10^{-8}$ resulting in a validation fit score of 0.92.

Table 3
Best-fit Sensitivity to Number of Shuffle Splits

# Splits	Best γ Trial 1	Best γ Trial 2	Best γ Trial 3
7	2.51×10^{-5}	3.98×10^{-6}	2.51×10^{-5}
10	2.51×10^{-5}	2.51×10^{-5}	2.51×10^{-5}
15	3.98×10^{-6}	2.51×10^{-5}	3.98×10^{-6}
20	3.98×10^{-6}	3.98×10^{-6}	3.98×10^{-6}
150	3.98×10^{-6}	3.98×10^{-6}	3.98×10^{-6}

Figure 8. Example depth determination for a 63.0 nm light curve.

polynomial fits: the larger the value, the more complex the model. We test γ values between 10^{-10} and 10^{-5} with 20 equal logarithmic steps in that range.

A major paradigm in machine learning is splitting data into training and validation sets. For time-series data such as this, it would not be appropriate to select, say, the first half of the points as the training set and the second half as the validation set because this could easily fail to capture important variability in the data, such as the dimming profile. Instead, we randomly choose half of the points across the entire time range for training and the other half for validation. Random errors could be introduced in this process due to the short timescale variability, which is considered noise for our purposes. To mitigate the introduction of these errors, *scikit-learn*'s *ShuffleSplit* function has an optional input for the number of equivalent splits to do, i.e., more than one random selection of the testing and validation sets are produced. We performed a simple sensitivity test to determine how many splits were required to eliminate major differences in the best-fit results and found that 20 splits were sufficient (see Table 3).

To find the best γ, another machine-learning technique was applied: validation curves (Figure 6). As the complexity of the model increases (moving to the right in Figure 6), the explained variance score of the fit always increases for the training set. However, the score of the fit applied to the validation set has a peak somewhere in the middle, which is the best fit. For very low model complexities, the model does a poor job fitting either set, which, in traditional nomenclature, is referred to as under-fitting. At high model complexities, a large disparity appears where the score for the training set is high but is low for the validation set. In traditional terms, this is known as over-fitting, and in machine-learning language, this behavior is sometimes referred to as the model "memorizing" the training set (as opposed to actually "learning" the trend behind the
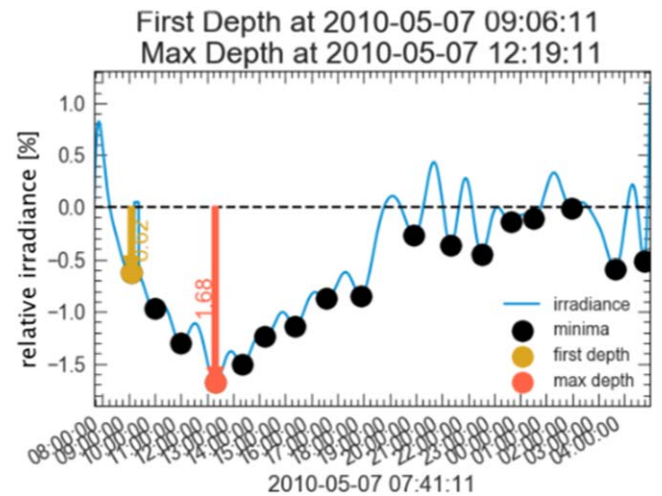
data). It is important to input a range of γ values that fully capture the validation curve so that a valid peak can be found. This process was completed for thousands of light curves, and their best-fit γ values were compared. We found that fits with $\gamma = 5 \times 10^{-8}$ always performed very well. Inspection of the fit light curve itself overlaid on the data (e.g., Figure 7) confirmed the desired behavior. Thus, to improve processing time for further runs of the catalog, we simply fix the value of γ. All subsequent steps in the algorithm use the light-curve fits rather than the underlying data.

3.4. Dimming Depth Determination

For each light-curve fit, the local minima are extracted using the *scipy.signal.argrelmin* function (Figure 8). Any minima above an irradiance value of 0% are thrown out because by definition, a dimming is a decrease from a baseline level. That baseline determination step (Section 3.1) normalizes the light curves so that 0% approximates the irradiance level that would be if no eruptive event and corresponding dimming occurred. If no minima remain after this filter is applied, then the dimming depth value is set to NaN. Otherwise, the earliest minimum defines the catalog's "depth first" and the deepest minimum defines the "depth max," which can be the same value.
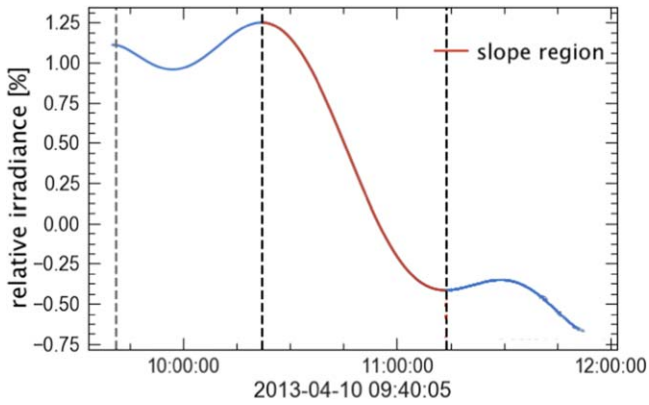
Figure 9. Example slope determination for a 13.1 nm light curve. The gray vertical dashed line indicates the *GOES*/XRS flare peak time, and the two black dashed lines are the edges of the slope window. The minimum slope value is $3.5 \times 10^{-6}\%\,\mathrm{s}^{-1}$, maximum is $9.1 \times 10^{-4}\%\,\mathrm{s}^{-1}$, and mean is $5.3 \times 10^{-4}\%\,\mathrm{s}^{-1}$.

3.5. Dimming Slope Determination

The determination of slope for a particular light curve depends on the successful determination of dimming depth (Section 3.4), which is used to constrain the end time of the window used to calculate the slope. Having already found the earliest local minimum slightly simplifies the algorithm at the cost of introducing a dependency. The start time of the slope window is initially defined to be the time of the associated *GOES*/XRS flare peak but is then refined to be the time of the maximum value in that window. This is because the *GOES*/XRS flare peak is representative of plasma as hot or hotter than the line emission in *SDO*/EVE and as such the flare peaks tend to be at later times in EVE than in XRS as the emitting plasma cools. The slope window should not include the flare peak or slightly earlier times because this would bias the slope value downward. Once this window is defined, the derivative is taken across the time range. The minimum, maximum, and mean derivatives are then recorded as slope values (Figure 9).

3.6. Dimming Duration Determination

The duration here is simply defined as the time between the first drop below 0% irradiance and the first time the irradiance rises back above 0% (Figure 10). Again, we assume that the baseline determination step (Section 3.1) has already normalized the light curve such that 0% is the baseline. The code first finds all zero crossings. It then discards any that occur before the slope start time for consistency of these two results. If no zero crossings remain, the code sets the duration to NaN. Otherwise, the code finds the first zero crossing that has a negative slope and the first subsequent zero crossing with a positive slope. In this way, we ensure that the duration we calculate corresponds to a depression in the light curve rather than an enhancement. The duration is then simply defined as the time between these two identified zero crossings.

4. JEDI Catalog Statistics

The catalog itself is a sparse array despite its large size. It is 5051 rows by 27,349 columns, where 90% of these 138,139,799 cells are NaN. Section 4.1 explains why. This still leaves 14,227,655 dimming and metadata parameters to work with. These are characterized in Section 4.2.
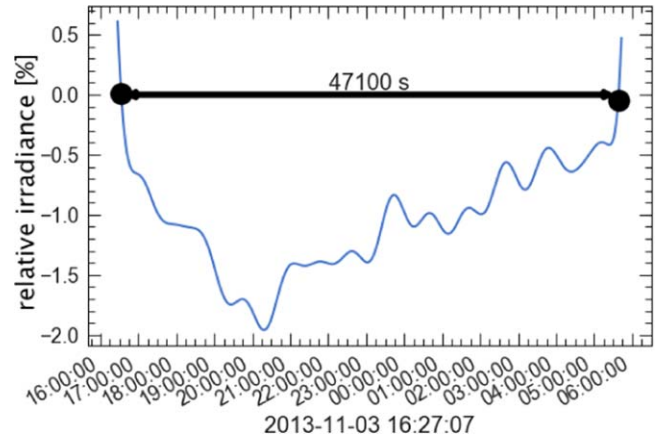

Figure 10. Example duration determination for a 25.6 nm light curve.

4.1. Algorithm Throughput

Each step of the algorithm can result in some light curves being rejected. In these cases, the JEDI catalog contains an NaN value. In the very first step, pre-flare irradiance determination, only 30% of light curves make it through (see Section 3.1.1).

The second step determines the size of the dimming time window, which can be truncated by subsequent flare interruptions, and has a throughput of 56%. The rejected light curves are those that have a subsequent flare less than 2 hr later. Note that the percentage of flares that have a subsequent flare less than 24 hr later is 96%, but the algorithm will still attempt to characterize those using a window truncated to the time of the subsequent flare.

The next step is the peak-match-subtract within the dimming window, which has a throughput of 14%. There are several reasons that the algorithm rejects the light curves. Fifty-one percent of the failures are due to a mismatch in the length of the two light curves (one to subtract with and one to subtract from). This is so common because MEGS-A and MEGS-B often have dramatically different effective cadences of observation due to the radiation issues that impacted MEGS-B early in the mission (Section 2.1). Next, 34% of the failures are due to excessive time shifts. If the peaks identified in the two light curves are more than 30 minutes apart, the algorithm rejects them. Finally, the algorithm also rejects the light curves if a peak cannot be identified in either light curve, which accounts for 15% of the failures.

For each potential event, we then fit all of the light curves: those corresponding to the original 39 emission lines and every one of their 1482 permutations from the previous step. The throughput of this fitting is 24%. The failed cases are those whose best-fit scores were less than our threshold value of 0.5.

Finally, the fitted light curves are parameterized in terms of depth, slope, and duration, if possible. The depth throughput is 80%. It fails if either (1) no local minima are found or (2) none of the minima are below the pre-flare irradiance baseline. Note that the slope determination also depends on the depth being identified, because it uses the time of the first minimum as the end time for slope determination. In light curves where that time could be provided, the slope algorithm throughput is 100%. The duration determination has a throughput of 74%. Four conditions can cause it to reject the light curve: 40% of the duration failures are because the light curve never crosses the pre-flare irradiance baseline; 13% because no negative-slope baseline crossing can be found (i.e., the light curve never drops below the baseline); 47% because no positive-slope
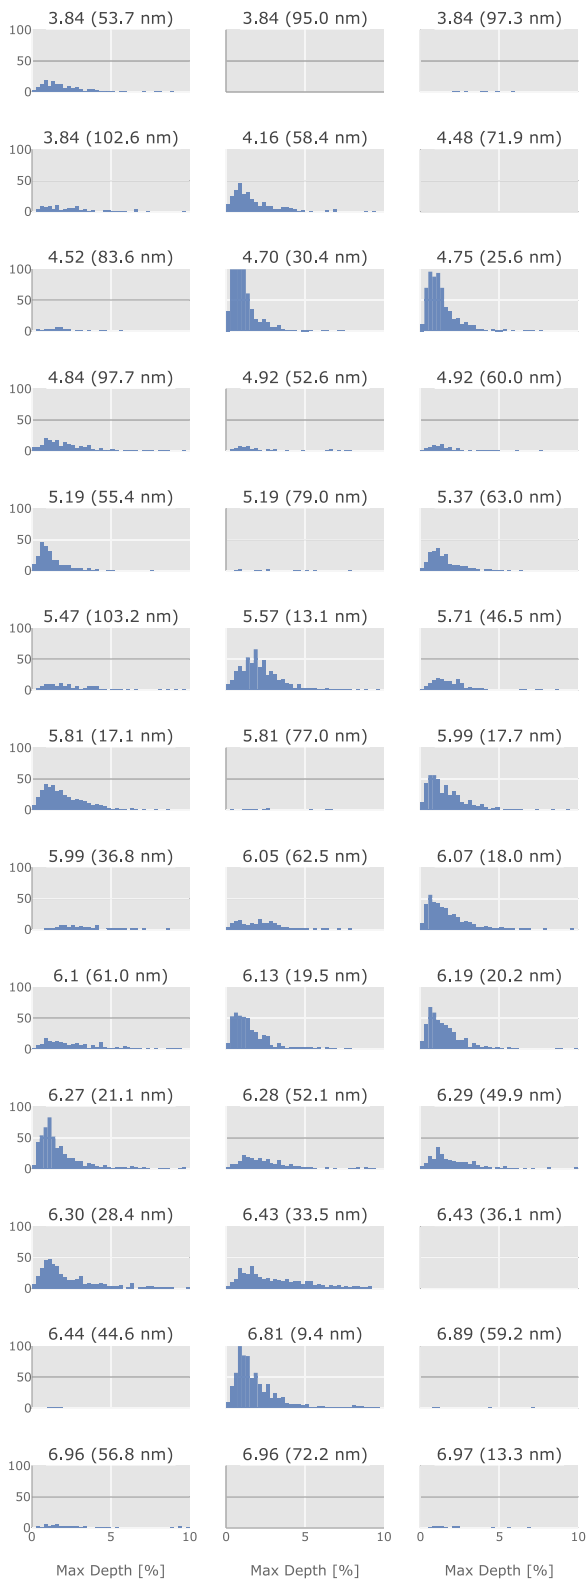
Figure 11. Dimming depth histograms sorted by the log of temperature in kelvin indicated in each plot title with its corresponding spectral line wavelength. The *x* and *y* axes are the same for all plots.

Figure 12. Same as Figure 11 but normalized by instrument duty cycle. Note that the *y*-range is doubled.

baseline crossing can be found (i.e., the light curve never rises back up to the baseline), and 0% because the positive-slope crossing occurs before the negative-slope crossing.

Combining the above throughputs, we get $30\% \times 56\% \times 14\% \times 24\% = 0.56\%$ before parameterization. Depth and

slope then are 0.45%. Duration, which is independent of the other two parameters, is 0.42%. Despite these small percentages, they are applied to several million light curves. The catalog contains 16,642 depths, 16,577 slopes, and 12,333 durations.
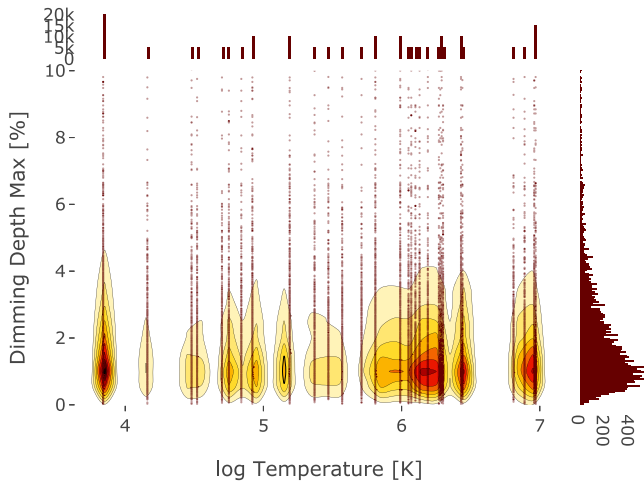
9

Figure 13. 2D density plot of dimming depth and formation temperature for the 39 emission lines, normalized by instrument duty cycle. The raw data are represented by the points and are overlaid by the contoured heat map, which indicates the point density. The histograms on the top and right indicate the density in their corresponding axis.

4.2. Histograms and Heat Maps

Figure 11 shows histograms of the maximum depths for the 39 uncorrected emission lines; the 1482 combined permutations of these lines are not shown here due to space limitations. The log temperature 4.70 (30.4 nm) histogram is cut off in order to allow visibility of all other plots. The height of the three cutoff bars are 210, 211, and 137. The large number of events in 30.4 and 25.6 may be due to obscuration dimming as described in Mason et al. (2014). Note that the plots corresponding to wavelengths >35 nm are from MEGS-B, which has significantly reduced operational cadence and therefore many fewer detections. The duty cycles are 96.3% and 16.8% observation up-time for MEGS-A and MEGS-B, respectively. Normalizing the histograms by these duty cycles results in Figure 12.

Figure 13 provides an alternative method of viewing all of the dimming max depth data for the 39 emission lines. The major population is at low dimming depths, with the greatest concentration of detections occurring in a range around a log temperature of 6. Note that there may be some systematic bias here because there happens to be a higher density of emission lines with temperatures just above a log temperature of 6 that physically should behave similarly and thus result in similar dimming parameterizations. Note that this is consistent with the conclusions of Mason et al. (2014, 2016), that CME-induced mass-loss dimming should occur around the ambient coronal temperature. The other hot spots are below a log temperature of 4—the same peak identified from 30.4 nm in Figures 11 and 12—and at a log temperature of 6.8. This latter hot spot is due to the 9.4 nm line, which may be compromised by the much cooler line blends from Fe IX and Fe X (Lepson et al. 2002; Schmelz et al. 2013).

Figure 14 is the same as Figure 13 but for the mean slope dimming parameter. While the slopes are a little more tightly clustered, as can be determined by looking at the vertical histogram scale, the general trends are similar to those of Figure 13.

5. Comparison to Prior Work

It is possible to compare some of the results from JEDI with prior irradiance coronal dimming work. Mason et al. (2014) performed a case study of a dimming event that occurred on 2010 August 7. JEDI rejected many of the light curves for this event due to significant oscillations in the pre-flare irradiance level, precluding a baseline determination by the algorithm described in Section 3.1. This oscillation, on the same timescale as the 120 minute window, caused the median comparison to exceed the threshold, as designed (Figure 15). However, one light curve studied by Mason et al. (2014) did get fully characterized by the JEDI algorithms: 19.5 nm (Figure 16). Surprisingly, the depth from JEDI is most similar to the corrected depths from Mason et al. (2014): 2.59% and 2.46%, respectively. The uncorrected depth from that paper, most comparable to Figure 16, was 1.52%. The difference likely stems from the pre-flare irradiance baseline because the other component of the depth value—identification of the minima—is simple. We believe the present work to be more reliable because the method for pre-flare irradiance baseline determination in Mason et al. (2014) was not established in as rigorous a manner as the algorithm described here; it only used a single irradiance measurement taken at 17:00 to define the baseline level. Note that in the case of the 17.1 nm emission line, this would correspond closely with the relatively high median value in the rightmost sub-window of Figure 15 that caused the present baseline determination algorithm to reject the light curve. This underscores the dependence on baseline determination techniques and the requirement that the same method be applied whenever comparing dimming depths.

We can also compare with the 29 irradiance dimming events in Mason et al. (2016). That study reported the dimming depth of 17.1 nm corrected by 28.4 nm, but despite the large size of this catalog, only 15 cases exist where both of those emission lines were simultaneously valid. However, an approximate comparison can be made. The mean and standard deviation of the 29 dimming depths in Mason et al. (2016) are 2.06% and 1.35%, respectively. In JEDI, for the 17.1 nm line—of which there are 402 values— they are 2.06% and 1.47%, respectively. This is an exceptional agreement, with the caveat that it is not a direct comparison.

6. Discussion and Future Work

The JEDI catalog represents a first step to systematically detect and characterize coronal dimmings in irradiance light curves. Dimming profiles contain information about the process that caused them, which is often due to a CME temporarily depleting a region of the coronal plasma. Mason et al. (2014, 2016) suggested that the depth of the dimming is indicative of CME mass and the slope of the dimming correlates with CME speed. This connection is especially important for exoplanet habitability studies because the most promising technique for detecting and characterizing stellar CMEs is irradiance coronal dimming parameterization (Harra et al. 2016). Furthermore, these methods may be useful for the development of new, simple instruments for solar space weather. Low-cost, low-mass, low-power, low-data irradiance instruments can be tuned to a few of the most important emission lines for determining CME speed and mass, as determined with the JEDI catalog in combination with the CME catalogs.

The natural next step for this work is to do precisely that: combine JEDI with existing large CME catalogs such as CDAW (Gopalswamy et al. 2009) and CACTus (Berghmans et al. 2002). These catalogs can be supplemented with input from CME experts and the *STEREO* coronagraphs, as was done in Mason et al. (2016). By doing so, we will be able to build upon that work to create a robust empirical relationship
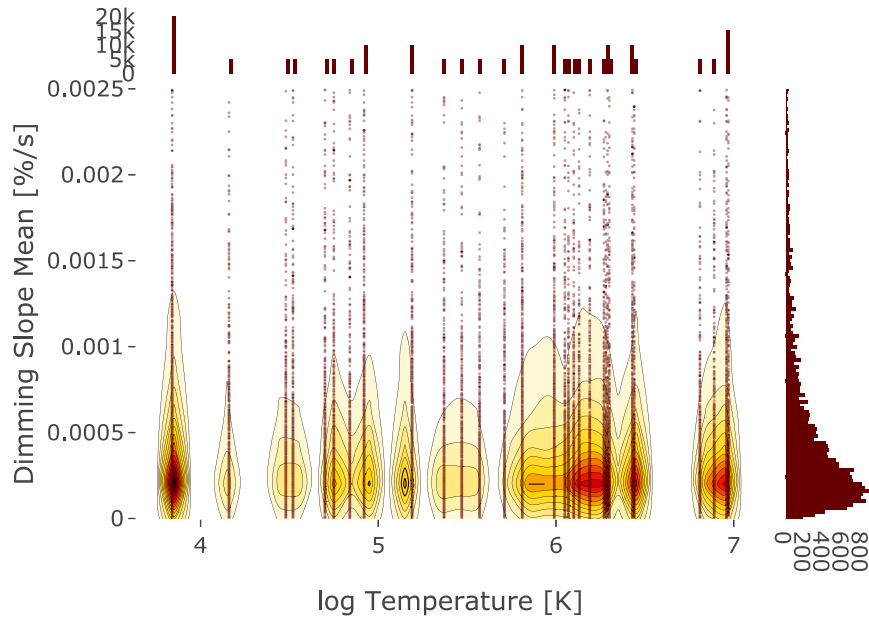
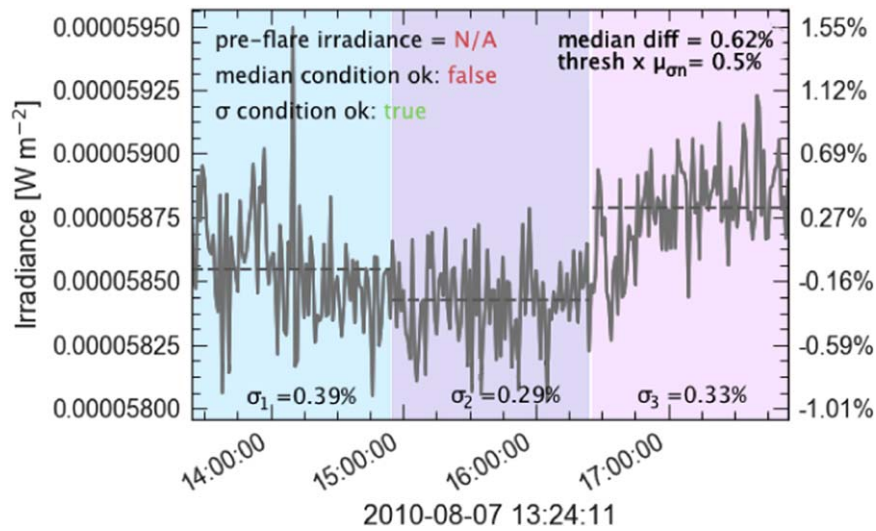**Figure 14.** Same as 13 but for the mean slope dimming parameter.

Figure 15. Baseline determination for the 17.1 nm line for the 2010 August 7 event that was also studied by Mason et al. (2014). While the σ condition passes, the median condition fails due to the irradiance increase in the rightmost sub-window.

between irradiance coronal dimming and CME kinematics. Theoretical work is being done in parallel to bolster this empirical relationship and further explore the physical connections (e.g., Jin et al. 2017). The results of these magnetohydrodynamic (MHD) simulations have so far confirmed that there is indeed a strong relationship, and they agree with the empirical relationship derived in Mason et al. (2016). Similarly, comparisons with image-based dimming studies have so far agreed with Mason et al. (2016; e.g., Dissauer et al. 2018a, 2018b). Meng Jin is presently running new MHD simulations with M-dwarf star parameters—stronger magnetic fields yielding significantly hotter coronae, and more powerful eruptions—to determine how the relationship between dimming and CMEs may differ from the solar case. No past or current astrophysical observations are capable of detecting such

dimmings, but these simulations can guide instrument requirements for potential future missions.

An additional next step is to compare JEDI with other dimming catalogs, for example, the Solar DEMON catalog (Kraaikamp & Verbeeck 2015) and the catalog produced by CoDiT (Krista & Reinard 2017). We can determine the percentage of mutual dimming detections and potentially use the other catalogs to guide thresholds to be applied to the parameters in JEDI for declaring whether or not each event should be considered a dimming.

Recall that this iteration of the JEDI catalog uses solar flares as the trigger to search for dimming. This means that non-flare-associated CMEs and dimmings are by definition excluded. Such CMEs tend to be less energetic (e.g., Howard & Harrison 2013; D'Huys et al. 2014), and as such, we expect
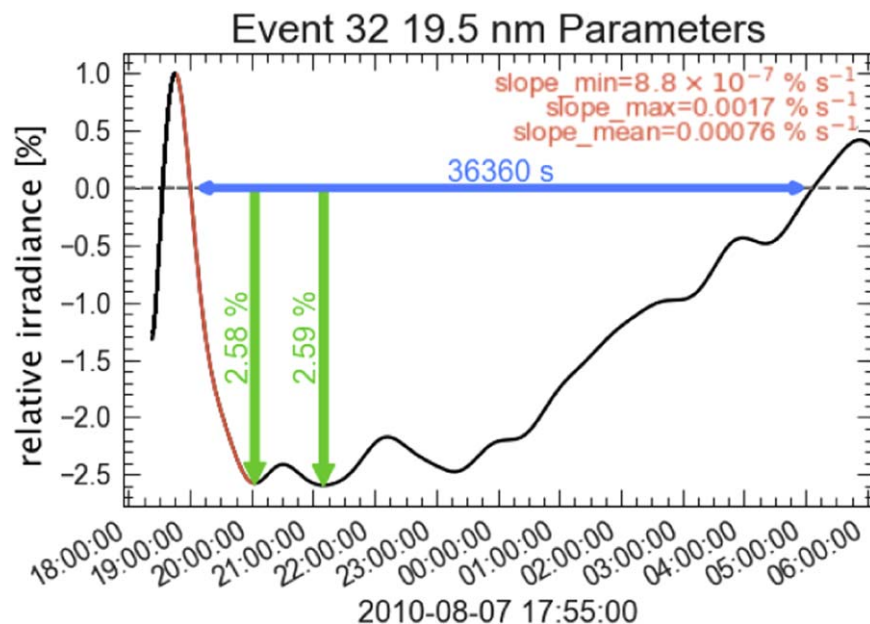
Figure 16. Summary plot for the 2010 August 7 event that was also studied by Mason et al. (2014).

that the irradiance dimming signals would be correspondingly degraded. In these CMEs, it may also be the case that space and temperature are less strongly coupled: heating may be insufficient to cause increased emission in the hotter lines but still cause increases in the cooler emissions that could partially or totally mask irradiance dimming. The present work has started with the easier problem of flare-associated CMEs and dimming. However, we are already beginning development of an irradiance dimming catalog that uses the Solar Demon dimming catalog to trigger searches, rather than the *GOES* flare event list. This new catalog is tentatively titled SITH.

Finally, the code to generate JEDI and the catalog itself are entirely open source and version controlled. Future improvements will continue to be documented, and we encourage both community input directly on the code repository and community usage of the results.

ORCID iDs

James Paul Mason ⓘ https://orcid.org/0000-0002-3783-5509
Raphael Attie ⓘ https://orcid.org/0000-0003-4312-6298
Barbara Thompson ⓘ https://orcid.org/0000-0001-6952-7343
Thomas N. Woods ⓘ https://orcid.org/0000-0002-2308-6797

References

AAS Journals Team, & Hendrickson, A. 2018, AASJournals/AASTeX60, v6.2, Zenodo, doi:10.5281/ZENODO.1209290
Aschwanden, M. J. 2009, AnGeo, 27, 3275
Berghmans, D., Foing, B. H., & Fleck, B. 2002, ESA Spec. Publ., 508, 437
D'Huys, E., Seaton, D. B., Poedts, S., & Berghmans, D. 2014, ApJ, 795, 49
Dissauer, K., Veronig, A. M., Temmer, M., & Podladchikova, T. 2018a, ApJ, 874, 123
Dissauer, K., Veronig, A. M., Temmer, M., Podladchikova, T., & Vanninathan, K. 2018b, ApJ, 863, 169
Drucker, H., Burges, C. J. C., Kaufman, L., Smola, A., & Vapoik, V. 1997, Support Vector Regression Machines, Tech. Rep., (New Providence, NJ: Bell Labs)
Gopalswamy, N., Yashiro, S., Michalek, G., et al. 2009, EM&P, 104, 295
Harra, L. K., Schrijver, C. J., Janvier, M., et al. 2016, SoPh, 291, 1761
Howard, T. A., & Harrison, R. A 2013, SoPh, 285, 269
Hudson, H. S., Acton, L. W., & Freeland, S. L. 1996, ApJ, 470, 629
Hunter, J. D. 2007, CSE, 9, 90
Jin, M., Cheung, C. M. M., DeRosa, M. L., Nitta, N., & Schrijver, K. 2017, AGUFM, SH41A-2758
Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python v3.5.0, http://www.scipy.org/
Kahler, S. W. 1982, JGR, 87, 3439
Kraaikamp, E., & Verbeeck, C. 2015, JSWSC, 5, 18
Krista, L. D., & Reinard, A. A. 2017, ApJ, 839, 50
Lemen, J. R., Title, A. M., Akin, D. J., et al. 2012, SoPh, 275, 17
Lepson, J. K., Beiersdorfer, P., Brown, G. V., et al. 2002, ApJ, 578, 648
Ma, S., Attrill, G. D. R., Golub, L., & Lin, J. 2010, ApJ, 722, 289
Mason, J. P. 2018, Python Convenience Functions v1.0.0, Zenodo, doi:10.5281/zenodo.1463150
Mason, J. P. 2019, James's EVE Dimming Index (JEDI) v2.0.0, Zenodo, doi:10.5281/zenodo.3243634
Mason, J. P., Woods, T. N., Caspi, A., Thompson, B. J., & Hock, R. A. 2014, ApJ, 789, 61
Mason, J. P., Woods, T. N., Webb, D. F., et al. 2016, ApJ, 830, 12
McKinney, W. 2010, in Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman (Austin, TX: SciPy), 51, https://conference.scipy.org/proceedings/scipy2010/pdfs/mckinney.pdf
Oliphant, T. E. 2006, Guide to NumPy
Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, J. Mach. Learn. Res., 12, 2825, http://jmlr.org/papers/v12/pedregosa11a.html
Perez, F., & Granger, B. E. 2007, CSE, 9, 21
Pesnell, W. D., Thompson, B. J., & Chamberlin, P. C. 2012, SoPh, 275, 3
Plotly Technologies Inc. 2015, Collaborative Data Science (Montreal, QC: Plotly Technologies Inc.)

Mason et al.

Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, AJ, 156, 123
Reinard, A. A., & Biesecker, D. A. 2008, ApJ, 674, 576
Rust, D. M., & Hildner, E. 1976, SoPh, 48, 381
Ryan, D. F., Milligan, R. O., Gallagher, P. T., et al. 2012, ApJS, 202, 15
Schmelz, J. T., Jenkins, B. S., & Kimble, J. A. 2013, SoPh, 283, 325
Thompson, B. J., Cliver, E. W., Nitta, N. V., Delannée, C., & Delaboudiniere, J. P. 2000, GeoRL, 27, 1431
Thompson, W. T. 2006, A&A, 449, 791

Waskom, M., Botvinnik, O., O'Kane, D., et al. 2017, mwaskom/seaborn: v0.8.1, Zenodo, doi:10.5281/ZENODO.883859
Woods, T. N., Eparvier, F. G., Hock, R. A., et al. 2010, SoPh, 275, 115
Woods, T. N., Hock, R. A., Eparvier, F. G., et al. 2011, ApJ, 739, 59
Yashiro, S., Gopalswamy, N., Akiyama, S., Michalek, G., & Howard, R. A. 2005, JGR, 110, A12S05
Zarro, D. M., Sterling, A. C., Thompson, B. J., Hudson, H. S., & Nitta, N. V. 1999, ApJL, 520, L139