
INCORPORATING MAGNETIC FIELD CHARACTERISTICS INTO EUV-BASED AUTOMATED SEGMENTATION OF CORONAL HOLES

Jeremy A. Grajeda, Laura E. Boucheron

Klipsch School of Electrical and Computer Engineering
New Mexico State University
Las Cruces, NM 88003, USA
{jgra, lboucher}@nmsu.edu

Michael S. Kirk

Heliophysics Space Division, Goddard Space Flight Center
National Aeronautics and Space Administration
Greenbelt, MD 20771, USA
michael.s.kirk@nasa.gov

Andrew Leisner

Department of Physics and Astronomy
George Mason University
Fairfax, VA 22030, USA
aleisner@gmu.edu

C. Nick Arge

Heliophysics Space Division, Goddard Space Flight Center
National Aeronautics and Space Administration
Greenbelt, MD 20771, USA
charles.n.arge@nasa.gov

Jaime A. Landeros

Heliophysics Space Division, Goddard Space Flight Center
National Aeronautics and Space Administration
Greenbelt, MD 20771, USA
University of California San Diego
La Jolla, CA 92093, USA
ADNET Systems Inc
Bethesda, MD 20817, USA
jalanderos@ucsd.edu

ABSTRACT

Coronal holes (CH) are magnetically open regions that allow hot coronal plasma to escape from the Sun and form the high-speed solar wind. This wind can interact with Earth's magnetic field. For this reason, developing an accurate understanding of CH regions is vital for understanding space weather and its effects on Earth. The process of identifying CH regions typically relies on extreme ultraviolet (EUV) imagery, leveraging the fact that CHs appear dark at these wavelengths. Accurate identification of CHs in EUV, however, can be difficult due to a variety of factors, including stray light from nearby regions, limb brightening, and the presence of filaments (which also appear dark, but are not sources of solar wind). In order to overcome these issues, this work incorporates photospheric magnetic field data into a classical EUV-based segmentation algorithm based on the active contours without edges (ACWE) segmentation method. In this work magnetic field data are incorporated directly into the segmentation process, serving both as a method for removing non-CH regions in advance, and as a method to constrain evolution of the boundary which identifies the CH boundary. This reduces the presence of filaments while simultaneously allowing the segmentation to include CH regions that may be difficult to identify due to the aforementioned inconsistent intensities.

Keywords Coronal Holes, automated detection; Coronal holes, magnetic fields

1 Introduction

Coronal holes (CHs) are regions of open magnetic field lines [Altschuler et al., 1972]. These lines, which emerge from the photosphere and extend through the Sun’s corona, allow coronal plasma to escape from the Sun and form high speed flows of particles called the solar wind [Wang and Sheeley, Jr., 1990, Wang et al., 1996, Antonucci et al., 2004, McComas et al., 2007] that can affect Earth’s magnetic field [Tsurutani et al., 2006]. Due to this relationship between CHs and solar wind, accurate detection and segmentation of CHs proves vital for understanding solar wind behavior, and, by extension, improving space weather forecasting. This can be seen, for example, in the development of the Wang-Sheeley-Argé (WSA) model [Wang and Sheeley, Jr., 1990, Arge et al., 2003, 2004].

Due to the lower density and temperature of CH regions, which is a byproduct of the open magnetic field lines, CHs can be detected in extreme ultraviolet (EUV) imagery, where they appear as dark regions [Munro and Withbroe, 1972]. This has resulted in a large number of detection schemes which rely on EUV imagery, including those developed by Krista and Gallagher [2009], Verbeeck et al. [2014], Lowder et al. [2014], Caplan et al. [2016], Boucheron et al. [2016], Garton et al. [2018], Hamada et al. [2018], Grajeda et al. [2023]. Each of these methods is developed to address difficulties with the segmentation process, including overcoming inconsistent intensities caused by stray light from nearby regions, limb brightening, and instrument effects [Verbeeck et al., 2014, Caplan et al., 2016]. In addition to these difficulties, the presence of filaments, which also appear dark in EUV but are not sources of solar wind, can result in the misclassification of regions when using only EUV data [Krista and Gallagher, 2009, Reiss et al., 2023].

The open magnetic field lines present in CHs emerge from highly unipolar regions unlike magnetic fields in filaments which are bipolar in nature. This has resulted in several methods that validate CH regions after detection through the use of magnetic field data, removing non-CH regions as a post-hoc process in the pipeline [Scholl and Habbal, 2008, Krista and Gallagher, 2009, Lowder et al., 2014, Garton et al., 2018, Hamada et al., 2018, Landeros et al., 2024]. With that noted, direct application of magnetic field data during the initial detection and segmentation of CH regions remains fairly limited, appearing predominately in Jarolim et al. [2021] where it serves as an additional channel alongside seven EUV channels input to a network based on the U-Net architecture of Ronneberger et al. [2015]. The work presented here proposes the incorporation of magnetic field data into the classical segmentation method first employed by Boucheron et al. [2016] and extended in Grajeda et al. [2023] which is based on the Active Contours Without Edges (ACWE) segmentation process of Chan and Vese [2001]. The ACWE-based method has already been shown to effectively overcome the inconsistent intensity present in CH observations for the majority of cases and is shown to be robust across short timeframes where CH evolution is expected to be minimal [Grajeda et al., 2023]. However, like all other EUV-based methods, it can misidentify filament regions [Grajeda et al., 2023, Reiss et al., 2023]. Through the addition of magnetic field data in both seeding and iterating of the ACWE process, the method proposed here reduces filament contamination, first by eliminating many filaments outright and second by reducing the overall area of the filaments that cannot be outright removed. Additionally, incorporation of magnetic field in the seeding and iterating of ACWE better constrains the evolution to avoid including too much quiet Sun (QS) as CH.

The rest of this work is organized as follows: Section 2 provides an introduction to the ACWE segmentation, outlines how this method was adapted for CH detection, and introduces the dataset that will be used to validate this work. Section 3 outlines the process of incorporating magnetic field data into ACWE evolution and explores the resulting effects. Section 4 outlines the process of incorporating magnetic field data into the initial seeding to help remove filament contamination outright. The results from Sections 3 and 4 are used to develop a final pipeline, which is discussed and validated in Section 5. Conclusions and discussion of future work is provided in Section 6.

2 Background

2.1 Active Contours Without Edges (ACWE)

The detection method proposed here is an extension of the ACWE algorithm developed by Chan and Vese [2001] and first applied to CH detection in Boucheron et al. [2016]. At its most basic, ACWE is a two-step process designed to separate an image into foreground (in our case the CHs) and background (the remaining solar features and all off-disk areas). In the first step, one or more enclosed shapes, collectively called the contour (C), are defined. This contour acts as an initial guess, with regions inside the contour (C^+) serving as the foreground and regions outside of the contour (C^-) serving as the background. In the second step, the contour is refined by evaluating every pixel along the boundary of the contour to determine if it is more similar to the pixels within the foreground or the background and redrawing C accordingly. The second step is repeated until a stopping criterion, such as no evolution between iterations, is met.

In the standard formulation of ACWE, the process of determining whether a pixel along the contour boundary belongs to C^+ or C^- is decided using the ACWE energy functional:

$$F(m^+, m^-, C) = \mu \ell(C) + \lambda^{H+} \int_{C^+} |I(x, y) - m^+|^2 dx dy + \lambda^{H-} \int_{C^-} |I(x, y) - m^-|^2 dx dy, \quad (1)$$

where I is the image. This equation summarizes three ‘‘forces’’ that act on C , giving each force a user-defined weight that determines the relative importance in defining the final contour. The contour is manipulated in order to minimize this energy functional. The first force,

$$\mu \ell(C), \quad (2)$$

attempts to minimize the length of the contour $\ell(C)$, and uses the user-defined weight μ to determine its relative importance. The second force,

$$\lambda^{H+} \int_{C^+} |I(x, y) - m^+|^2 dx dy, \quad (3)$$

attempts to create a ‘‘homogeneous’’ foreground, where homogeneous is understood to mean a narrow range of intensities. This is done by comparing every pixel of $I(x, y)$ in C^+ , to the mean intensity of the foreground m^+ . This force is subject to the user-defined weight λ^{H+} . The final force,

$$\lambda^{H-} \int_{C^-} |I(x, y) - m^-|^2 dx dy, \quad (4)$$

mirrors the behavior of the second force (Equation 3), this time trying to create a homogeneous background by comparing every pixel in C^- to the mean intensity of the background m^- . It is subject to the user-defined weight λ^{H-} .

In both Boucheron et al. [2016] and Grajeda et al. [2023], the initial contour (seed) is defined using an intensity threshold applied to a 193 Å EUV image that has been corrected for limb brightening using the method from Verbeeck et al. [2014]. This threshold defines all on-disk pixels with an intensity $\leq \alpha \times m_{QS}$ (where α is a user-defined parameter [typically 0.3] and m_{QS} is an estimate of the mean intensity of the QS) as CHs. The contour is evolved using the same 193 Å EUV image, after masking off-disk areas from the algorithm, and is stopped when either no evolution has occurred between iterations or the only evolution that occurs consists of pixels along the boundary that oscillate between foreground and background.

2.2 Dataset and Metrics

For this work we will be utilizing the dataset we previously developed in Grajeda et al. [2023]. This dataset contains AIA Level 1 EUV images at 94, 131, 171, 193, 211, 304, and 335 Å, and the temporally corresponding 720-second line-of-sight (LOS) HMI magnetograms, collected at a one-hour cadence from Carrington rotations (CRs) 2099, 2100, 2101, and 2133. We note that all magnetogram data have an observation time within ± 2 s of the record time for the corresponding EUV data. We also note that LOS magnetograms may report inaccurate magnetic field at the solar limb, but that our method should be able to equivalently use radial magnetic field from vector magnetograms (an area of future research). Our workflow will be the reverse of the process in Grajeda et al. [2023], tuning our proposed variant of the ACWE algorithm to operate during CR 2133, which has a high level of filament activity, to ensure that this version of the algorithm can minimize false positive detections. Once the algorithm is tuned, the observations from CRs 2099-2101 will be used to verify that the algorithm continues to perform effectively during the periods of low solar activity where it previously excelled.

To allow for an accurate comparison between Grajeda et al. [2023] and this work, we will be utilizing the same four metrics: intersection over union (IOU), structural similarity index measure (SSIM), global consistency error (GCE), and local consistency error (LCE). Both IOU and SSIM report a value of 1 when two segmentations are identical, with IOU serving as a stringent method that reports the overlap of the two segmentations normalized by total area and SSIM providing a perceptual measure that is more consistent with the human visual system [Wang et al., 2004]. Both GCE and LCE report a value of 0 when two segmentations are identical. GCE and LCE are designed to report minimal or no error when one segmentation is a refinement of the other, with GCE serving as a more stringent metric [Martin et al., 2001].

3 Quantifying Unipolarity via Active Contour Kinetics

Incorporation of magnetic field data into ACWE evolution requires two adjustments to the standard formulation of ACWE. First, the ACWE algorithm must be altered to allow for vector-valued images so that EUV and HMI

magnetogram data can be simultaneously considered. Second, the ACWE energy functional must be altered to maximize unipolarity as expressed by the magnetic field data while still maximizing homogeneity in the EUV observation(s).

3.1 ACWE for Vector-Valued Images

The vector-valued formulation of ACWE was developed by Chan et al. [2000] by redefining the ACWE energy functional as

$$F(\bar{m}^+, \bar{m}^-, C) = \mu\ell(C) + \int_{C^+} \frac{1}{N} \sum_{i=1}^N \lambda_i^{\text{H}^+} |I_i(x, y) - m_i^+|^2 dx dy + \int_{C^-} \frac{1}{N} \sum_{i=1}^N \lambda_i^{\text{H}^-} |I_i(x, y) - m_i^-|^2 dx dy. \quad (5)$$

This new formulation evaluates the homogeneity of C^+ and C^- on a per-channel basis by comparing the pixels (foreground and background, respectively) in each channel I_i of the vector-valued image I to the channel-specific means m_i^+ and m_i^- . Within this formulation the weights for the two homogeneity parameters are also defined on a per-channel basis as $\lambda_i^{\text{H}^+}$ and $\lambda_i^{\text{H}^-}$ (foreground and background, respectively), allowing the user to adjust the relative influence each channel's foreground and background has on the final contour. It should be noted that when the total number of channels $N = 1$, the formulation in Equation 5 reduces to the original ACWE energy functional in Equation 1.

3.2 Incorporating Magnetic Unipolarity

As originally formulated, the vector-valued version of ACWE only considers the homogeneity of the underlying region within each channel. For this reason, the energy functional in Equation 5 must be modified to consider homogeneity in EUV only, while instead considering unipolarity in the corresponding magnetograms. Forcing ACWE to ignore homogeneity when evaluating the magnetic field can be accomplished by setting $\lambda_j^{\text{H}^+} = 0$ and $\lambda_j^{\text{H}^-} = 0$ where j is the index of the HMI magnetogram data within the vector-valued image I . From there, a metric that minimizes when a region becomes more unipolar can be added to the vector-valued energy functional (Equation 5).

We use the metric of Ko et al. [2014] as the basis for our measure of unipolarity. Ko et al. [2014] define unipolarity as:

$$U = \frac{\langle |B| \rangle - \langle B \rangle}{\langle |B| \rangle}, \quad (6)$$

where $\langle \cdot \rangle$ denotes mean, $|\cdot|$ denotes absolute value, and B is the magnetic field. This metric is bound to the range $[0, 1]$, where 0 indicates a purely unipolar region and 1 indicates a purely bipolar region. This ensures that increasing unipolarity will minimize the energy functional.

In order to incorporate Equation 6 into the vector-valued energy functional in Equation 5, the effect of moving a pixel at the contour boundary into or out of the foreground (CH) must be quantified. This means that both $\langle |B| \rangle$ and $\langle B \rangle$ must be calculated on a per-pixel basis. This process can be simplified by relying on the fact that the mean of a vector x is defined as

$$m_x = \frac{1}{N} \sum_{i=1}^N x(i) = \sum_{i=1}^N \frac{x(i)}{N}. \quad (7)$$

If the mean $m_{x_{N-1}}$ of a subset of $x(i)$ for $i = [1, N-1]$ has already been calculated, then the mean of the full ensemble for $i = [1, N]$ can be expressed as

$$m_{x_N} = \frac{x(N) + (N-1)m_{x_{N-1}}}{N}. \quad (8)$$

Using this formulation for the mean, the effect of a pixel $I(x, y)$ along C on the unipolarity of the foreground (in the i^{th} channel) can be expressed as

$$\int_{C^+} \frac{1}{N} \sum_{i=1}^N \lambda_i^{\text{U}^+} \frac{|I_i(x, y) + a_i^+ n^+| - |I_i(x, y) + m_i^+ n^+|}{|I_i(x, y) + a_i^+ n^+|} dx dy, \quad (9)$$

where a_i^+ is the absolute mean of the foreground (of the i^{th} channel), and n^+ is the number of pixels in the foreground. Note that we are using $\lambda_i^{\text{U}^+}$ to refer to the user-defined weight. Likewise, the unipolarity of the background can now be

expressed as

$$\int_{C^-} \frac{1}{N} \sum_{i=1}^N \lambda_i^{U^-} \frac{\left| \frac{|I_i(x,y)| + a_i^- n^-}{n^- + 1} - \frac{|I_i(x,y) + m_i^- n^-|}{n^- + 1} \right|}{\frac{|I_i(x,y)| + a_i^- n^-}{n^- + 1}} dx dy. \quad (10)$$

Adding Equation 9 and Equation 10 to the vector-valued energy functional (Equation 5) yields our new energy functional:

$$\begin{aligned} F(\bar{m}^+, \bar{m}^-, \bar{a}^+, \bar{a}^-, C) = & \mu \ell(C) + \int_{C^+} \frac{1}{N} \sum_{i=1}^N \lambda_i^{H^+} |I_i(x,y) - m_i^+|^2 dx dy \\ & + \int_{C^-} \frac{1}{N} \sum_{i=1}^N \lambda_i^{H^-} |I_i(x,y) - m_i^-|^2 dx dy \\ & + \int_{C^+} \frac{1}{N} \sum_{i=1}^N \lambda_i^{U^+} \frac{\left| \frac{|I_i(x,y)| + a_i^+ n^+}{n^+ + 1} - \frac{|I_i(x,y) + m_i^+ n^+|}{n^+ + 1} \right|}{\frac{|I_i(x,y)| + a_i^+ n^+}{n^+ + 1}} dx dy \\ & + \int_{C^-} \frac{1}{N} \sum_{i=1}^N \lambda_i^{U^-} \frac{\left| \frac{|I_i(x,y)| + a_i^- n^-}{n^- + 1} - \frac{|I_i(x,y) + m_i^- n^-|}{n^- + 1} \right|}{\frac{|I_i(x,y)| + a_i^- n^-}{n^- + 1}} dx dy. \end{aligned} \quad (11)$$

Here, note that there are foreground and background means for the EUV channel(s) and the HMI channel.

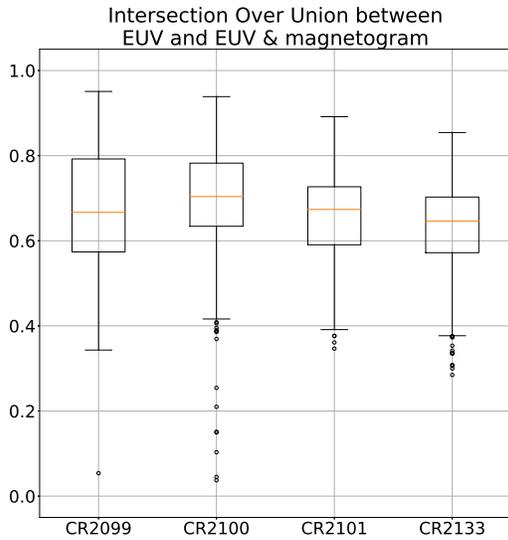
3.3 Effects of Magnetic Unipolarity on Segmentation

Following the pipeline established by Grajeda et al. [2023] for the creation of full-scale (4096×4096 pixels) CH segmentations, we begin by verifying that the header information of the Level 1 193 Å EUV image is correct (and updating it, if needed, using `aiapy.calibrate.update_pointing`), then converting to level 1.5 images using `aiapy.calibrate.register` [Barnes et al., 2020]. From there, we correct for limb brightening using the method of Verbeeck et al. [2014] and define the initial contour using the same threshold parameter ($\alpha = 0.3$). Magnetic field data are prepared by taking the corresponding HMI magnetogram and aligning it to the EUV observation using the reprojection tool, `reproject.reproject_interp`, of Robitaille et al. [2020]. The EUV and HMI data are combined into a two-channel image, with 193 Å EUV data in the first channel ($i = 1$) and HMI data in the last ($i = 2$). We note here that Ko et al. [2014] calculated unipolarity using an estimate of the radial magnetic field. The results herein presented are for the line-of-sight magnetic field, due to improved stability.

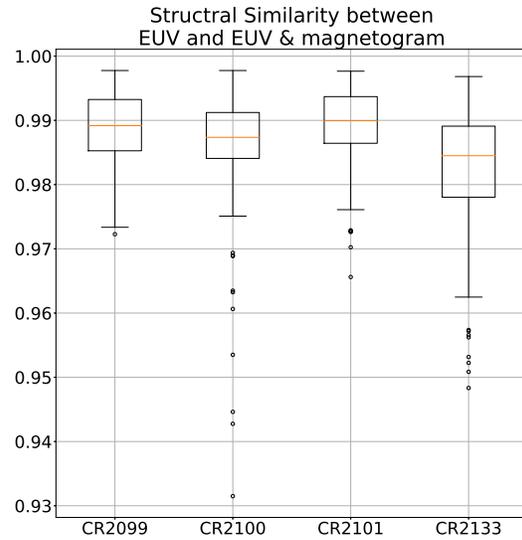
Matching Grajeda et al. [2023], off-disk areas in both channels are set to the mean intensity of the non-CH region. This operation is performed on a per-channel basis, and reset every 10 iterations. Evolution is performed using the same parameters as Grajeda et al. [2023] for contour length ($\mu = 0$), and for 193 Å homogeneity ($\lambda_1^{H^+} = 1$ and $\lambda_1^{H^-} = 1/50$). Additionally, for the EUV data, we set the unipolarity parameters $\lambda_1^{U^+} = \lambda_1^{U^-} = 0$. For the HMI data, we set the homogeneity parameters $\lambda_2^{H^+} = \lambda_2^{H^-} = 0$, and define the unipolarity parameters as $\lambda_2^{U^+} = 1$, and $\lambda_2^{U^-} = 1/50$.

The IOU, SSIM, GCE, and LCE between the full-scale 193 Å segmentations of Grajeda et al. [2023] and our EUV+HMI segmentations, organized by CR, are presented in Figure 1. We note that, across all four metrics, CR 2133 displays the largest median discrepancy. Given that filaments are most prevalent in CR 2133, this is consistent with the expectation that the magnetic unipolarity parameters are reducing filament contamination. Visual inspection of the segmentation pairs with the lowest IOU for CR2133, of which four samples are provided in Figure 2, confirms this speculation by showing that filament regions, while still present, are consistently diminished to a small region surrounding the initial seed.

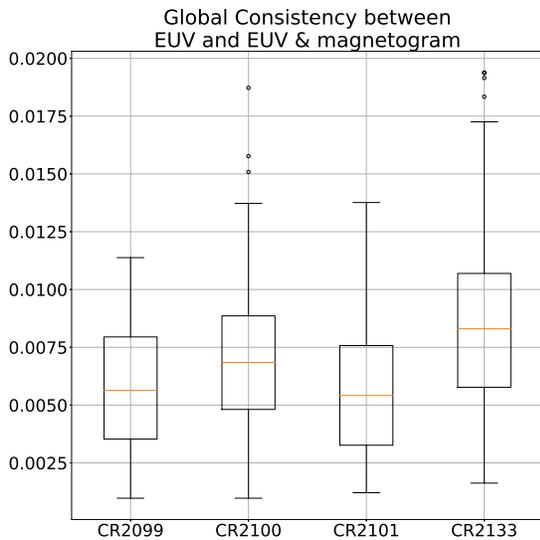
Referring back to Figure 1, we note that that CR 2099 shows a large range between Q1 and Q3 in IOU, GCE, and LCE, but not in SSIM. We also note that CR 2099 shows the second largest median discrepancy in IOU, but not in the remaining three metrics. Both of these results are consistent with a change in overall CH area while retaining the overall form or shape of the identified CHs. Grajeda et al. [2023] note that CR 2099 had the largest number of segmentations that changed from targeting CHs to QS as a result of including excessive QS in C^+ . This may suggest that our new formulation is reducing the growth of C^+ into QS regions, which is another motivation for inclusion of magnetic field into evolution rather than just post-hoc processing. Visual inspection of the segmentation pairs with the lowest IOU for CR 2099, of which four samples are provided in Figure 3, also confirms this speculation.



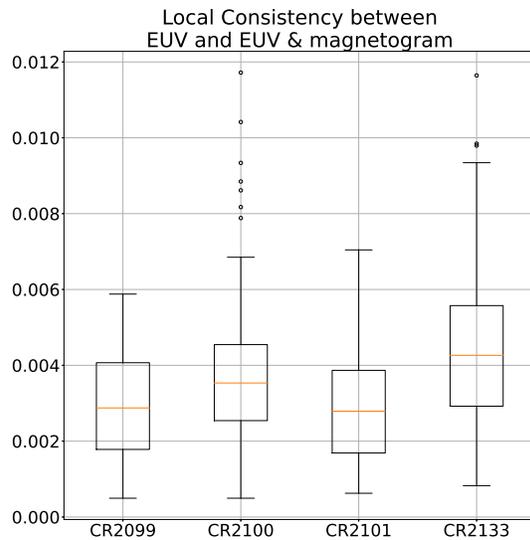
(a) Intersection Over Union



(b) Structural Similarity



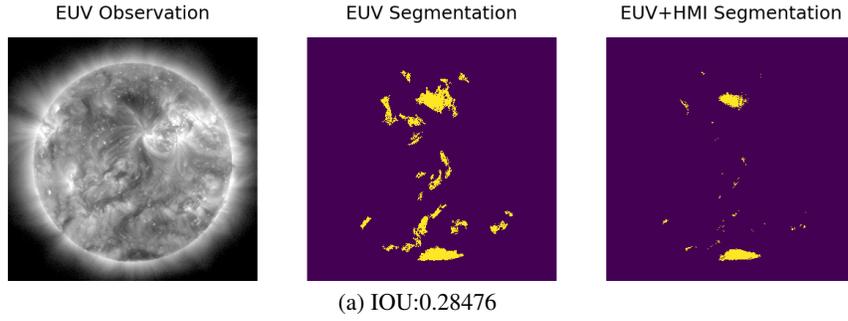
(c) Global Consistency Error



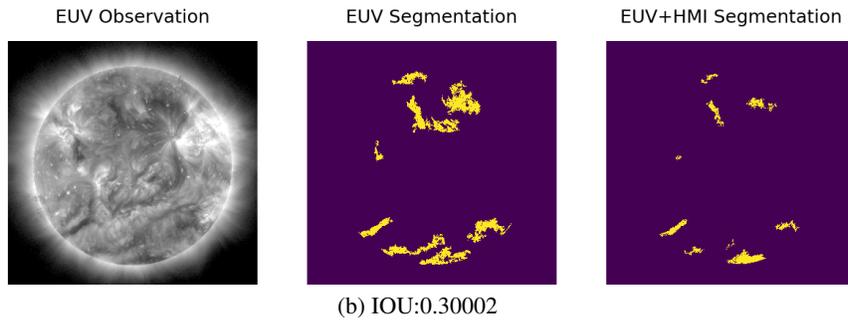
(d) Local Consistency Error

Figure 1: Comparison between full-scale segmentation results using only 193 Å data, and using both 193 Å and HMI magnetogram data, organized by CR. The box outlines the range between Q1 and Q3, with the median value in orange. The whiskers show 1.5 times the interquartile range. Outliers are marked with circles.

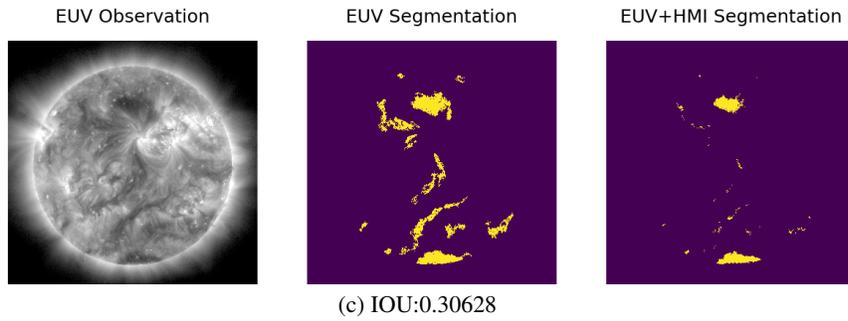
2013-02-12T01:00:01.000



2013-02-13T12:00:01.000



2013-02-11T19:00:01.000



2013-02-12T05:00:01.000

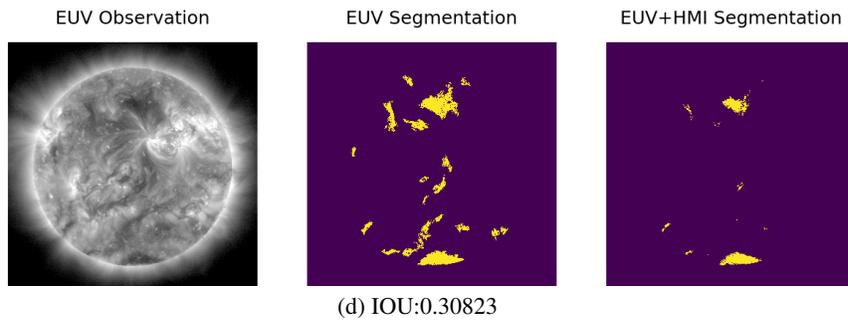
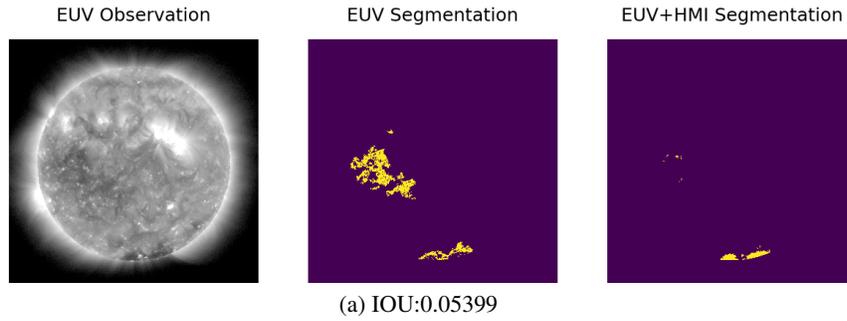
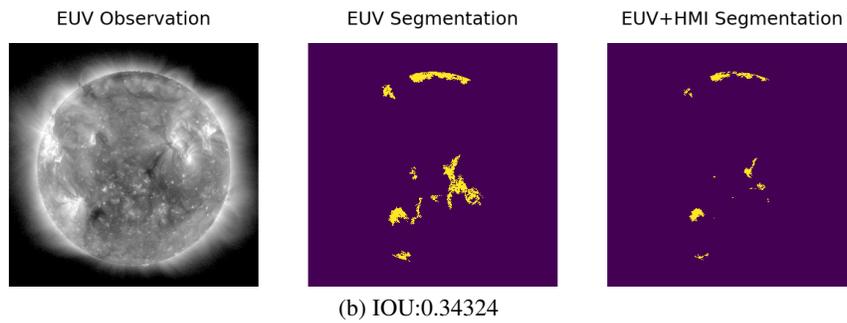


Figure 2: Example of full-scale(4096 × 4096 pixels) segmentation results using only 193 Å observations (center), and using both 193 Å and HMI magnetogram observations (right) over CR 2133. In each figure the leftmost image is the 193 Å observation. The title of each figure is the record time for the EUV data.

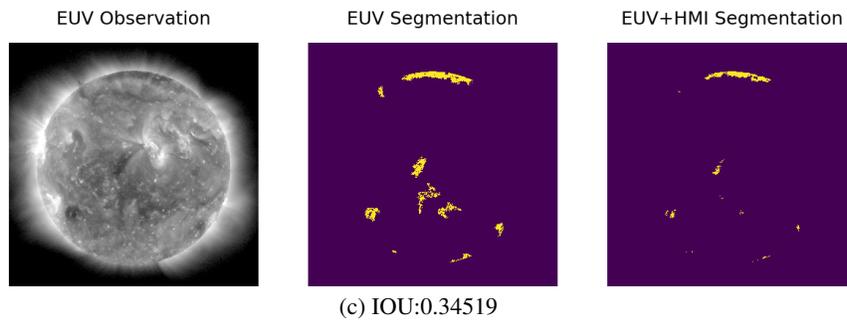
2010-07-16T23:00:02.000



2010-08-06T04:00:02.000



2010-08-04T20:00:02.000



2010-08-06T06:00:02.000

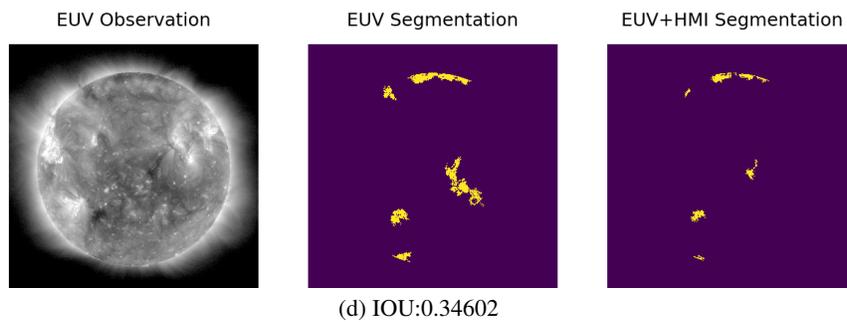


Figure 3: Example of full-scale (4096×4096 pixels) segmentation results using only 193 \AA observations (center), and using both 193 \AA and HMI magnetogram observations (right) over CR 2099. In each figure the leftmost image is the 193 \AA observation. The title of each figure is the record time for the EUV data.

In summary, incorporation of magnetic field information into ACWE (via a unipolarity metric operating on the HMI data) appears to have had negligible impact on the segmentation of CHs. At the same time, incorporation of the magnetic field appears to have minimized area of detected filaments and alleviated many cases of change of target. We will return to both of these results in Section 4 after first considering effects of spatial resolution on segmentation accuracy.

3.4 Effects of Spatial Resolution on Segmentation Accuracy

As noted in Grajeda et al. [2023], the iterative nature of ACWE results in segmentations at full-scale (4096×4096 pixels) taking several minutes to complete. Reducing the image to one-eighth-scale resolution (512×512 pixels) reduces segmentation time to a matter of seconds. In EUV-only CH segmentations, Grajeda et al. [2023] note that this process results in highly similar segmentations.

In order to determine the viability of using one-eighth-scale resolution images for EUV+HMI magnetogram segmentations, we replicate the pipeline of Grajeda et al. [2023] for producing segmentations at this reduced spatial resolution. This is done by reducing the scale of the 193 \AA image and the corresponding HMI magnetogram to 512×512 pixels using the scikit-image function `skimage.transform.resize` [van der Walt et al., 2014] after aligning the observations, but before correcting for limb brightening in EUV and before defining the seed using the EUV observation. We then evolve the initial seed on the decimated images, using the same weights as Section 3.3. After each segmentation has converged, it is upsampled back to the full 4096×4096 pixel resolution and compared with the EUV+HMI magnetogram segmentations developed in Section 3.3.

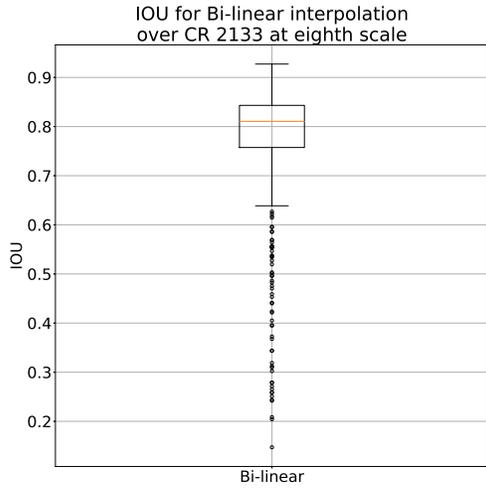
The effects of reducing spatial resolution on segmentation similarity for CR 2133 are summarized in Figure 4. Likewise, the effects of reducing spatial resolution on segmentation similarity for CRs 2099-2101 are summarized in Figure 5. In both figures, we provide similarity for bi-linear interpolation, noting that five additional interpolation methods (nearest-neighbor, bi-quadratic, bi-cubic, bi-quartic, and bi-quintic interpolation) were also tested and yielded similar results. We note here that the results in Figure 5 are consistent with the results in Grajeda et al. [2023], and that the results in Figure 4 are improved, suggesting that decimation to one-eighth-scale resolution remains a viable method for minimizing algorithm runtime despite (and in some cases, because of) the addition of decimated magnetogram data.

Visual inspection revealed the same three discrepancies between full- and reduced-resolution segmentations as in Grajeda et al. [2023]: absence of smaller regions due to the downsampling process removing the dark pixels that formed the initial seed, the presence or absence of spurious bright regions at different scales, and reduced fidelity along the contour boundary due to the lower spatial resolution. All three effects are demonstrated in Figure 6. No new discrepancies appear to have been added by introducing spatially downsampled HMI magnetogram data.

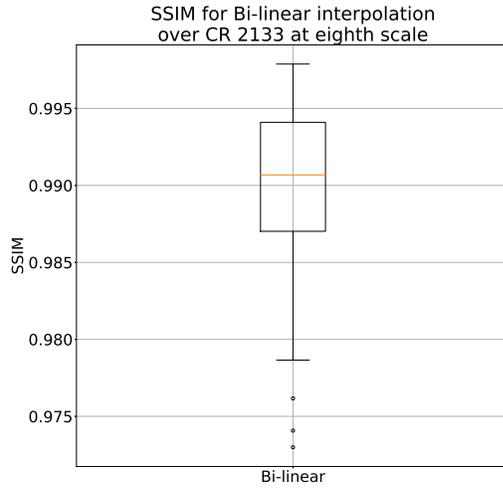
The first of the three effects, absence of smaller regions due to the downsampling process removing the dark pixels that formed the initial seed, is especially prevalent in the case of filaments, which, due to their thin structure, are more likely to lose dark pixels that form the initial seed when spatially downsampled. This effect explains the larger discrepancy that Grajeda et al. [2023] found in full- and reduced-scale segmentations for CR 2133. In particular, the fact that addition of HMI magnetogram data constrained filament evolution at the full-scale segmentation ensured higher similarity between the filament-absent one-eighth-scale segmentations, and the filament-present full-scale segmentations.

Due to the aforementioned absence of some filaments in CR 2133 at one-eighth-scale resolution, we additionally studied the effects of transferring the seed (initial contour) generated at full-scale resolution to the one-eighth scale resolution. This is done to ensure that all filaments present in the full-scale segmentations are also present in the one-eighth scale segmentations, and therefore to verify that the spatially-decimated HMI magnetogram data still proves effective at correctly constraining the CH segmentations. Due to differing spatial resolutions (4096×4096 pixels at full scale and 512×512 pixels at one-eighth scale), the full-scale seed must be downsampled to the new resolution. Since every pixel at the one-eighth-scale resolution represents an 8×8 pixel patch of the original image, in order to ensure that even a region represented with only one pixel in the full-scale seed is retained at the one-eighth-scale resolution, we apply dilation using a 4×4 pixel square using the scikit-image function `skimage.morphology.dilation` [van der Walt et al., 2014] prior to downsampling the initial seed. The effects of this process are summarized in Figure 7. We note that this process had minimal effect on SSIM, GCE, and LCE, but further improved IOU, suggesting that the added filaments were constrained in the same manner as before. This is confirmed through a visual inspection of the data. We provide, for reference, Figure 8, which provides the full-scale and one-eighth-scale segmentations generated utilizing the seed-transfer process for the same four examples from Figure 2. Transferring the seed for the remaining three CRs had no appreciable effect.

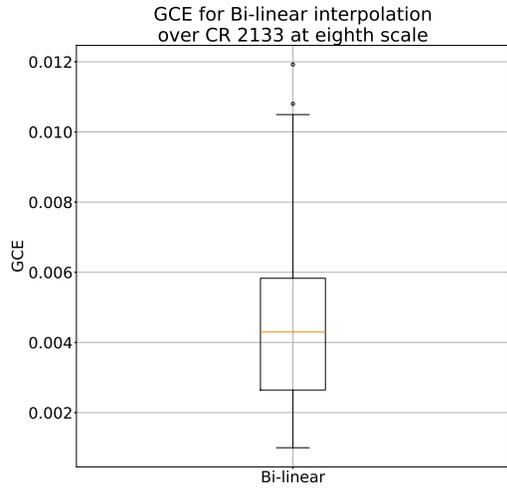
In summary, reducing the resolution of HMI magnetograms via standard image resizing procedures does not appear to affect the segmentation performance. We note that this image decimation does not take into account magnetic flux conservation (i.e., correction for differing areas of pixels across the disk), but appears to retain statistics relevant to measures of unipolarity. Additionally, we find that a loss of seeds (individual regions in the initial contour) related to



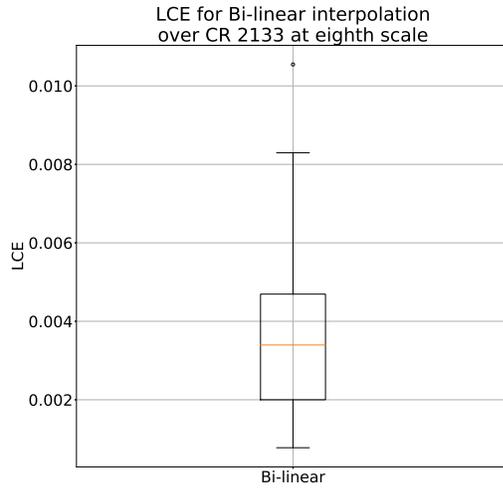
(a) Intersection Over Union



(b) Structural Similarity

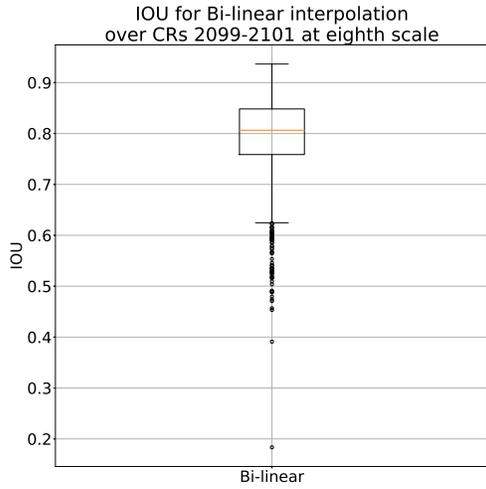


(c) Global Consistency Error

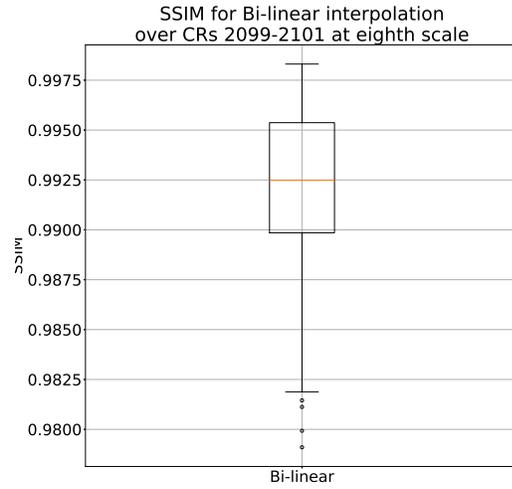


(d) Local Consistency Error

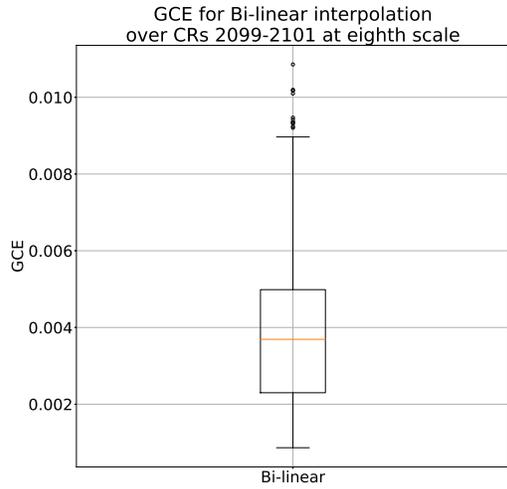
Figure 4: Comparison between full-scale (4096×4096 pixels) segmentation results using both 193 \AA and HMI magnetogram data, and one-eighth-scale (512×512 pixels) segmentation results using both 193 \AA and HMI magnetogram data for CR 2133. The box outlines the range between Q1 and Q3, with the median value in orange. The whiskers show 1.5 times the interquartile range. Outliers are marked with circles.



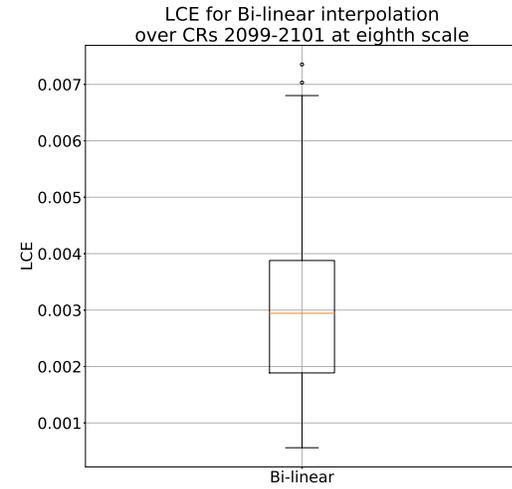
(a) Intersection Over Union



(b) Structural Similarity



(c) Global Consistency Error

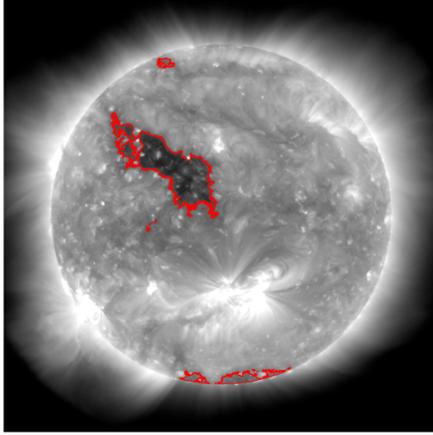


(d) Local Consistency Error

Figure 5: Comparison between full-scale (4096×4096 pixels) segmentation results using both 193 \AA and HMI magnetogram data, and one-eighth-scale (512×512 pixels) segmentation results using both 193 \AA and HMI magnetogram data for CRs 2099-2101. The box outlines the range between Q1 and Q3, with the median value in orange. The whiskers show 1.5 times the interquartile range. Outliers are marked with circles.

2010-09-18T00:00:02.000

Full Scale
(4096 × 4096 pixels)



Eighth Scale
(512 × 512 pixels)

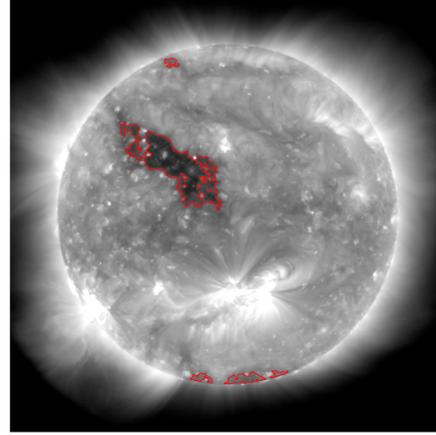


Figure 6: Example of full-scale (4096×4096 pixels) and one-eighth-scale (512×512 pixels) segmentation both generated using 193 \AA and HMI magnetogram data. The title of the figure is the record time of the EUV data.

filaments does not appear to be the largest contributing source of this performance. Rather, it appears that reduced-resolution HMI data retains relevant information for constraining evolution of the contour, reducing the contribution of filaments. Since filaments are still included in the segmentations, albeit with reduced area, we now consider incorporation of magnetic field information into the seeding process to further reduce filament contamination. This reinforces the importance of incorporating magnetic field information throughout the process of CH segmentation rather than as a post-hoc process.

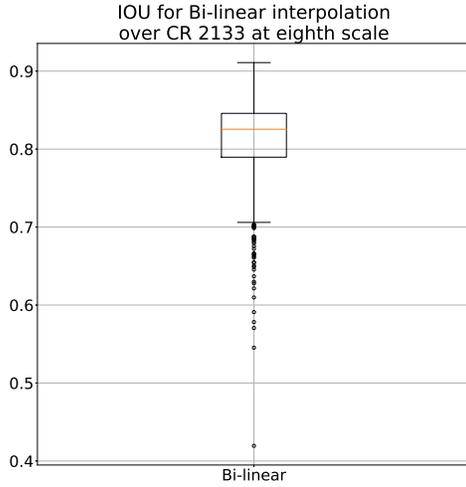
4 Incorporation of Magnetic Field Data into the Seeding Process

Due to improved runtime and minimal effect on segmentation quality, we use the one-eighth-scale segmentations for the second step (iterative refinement of the contour) in ACWE, even when utilizing HMI magnetogram data. Additionally, due to reduced filament contamination in the initial contours generated at one-eighth-scale resolution, we seed the algorithm after spatially downsampling the EUV and HMI magnetogram observations. With that noted, we recognize that while this formulation has improved the quality of the segmentations generated using ACWE, filament contamination, though minimized, persists. This section, therefore, explores the use of magnetic field data in initial seeding in order to further reduce filament contamination.

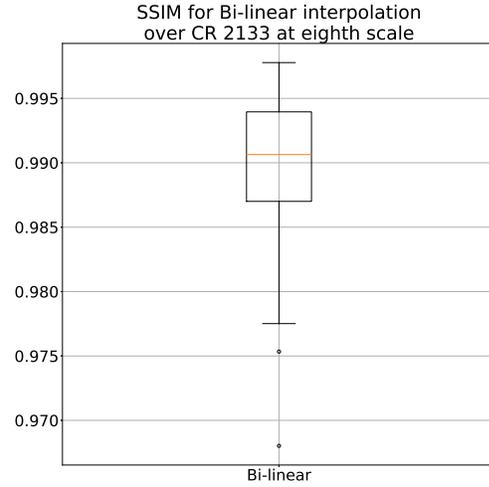
4.1 Complications with Estimating Unipolarity from Initial Seeds

Predicting the unipolarity of a final evolved region through direct evaluation of the seed, regardless of the scale at which seeding occurred, is not possible due to two issues. First, multiple seeds (individual regions in the seed/initial contour) often contribute to a single CH. Second, these seeds are often small regions, and are therefore vulnerable to the effects of small-scale statistics which results in erroneous estimations of unipolarity at the seeding stage. In order to demonstrate this, we will briefly return to the EUV-only segmentations for the following two studies.

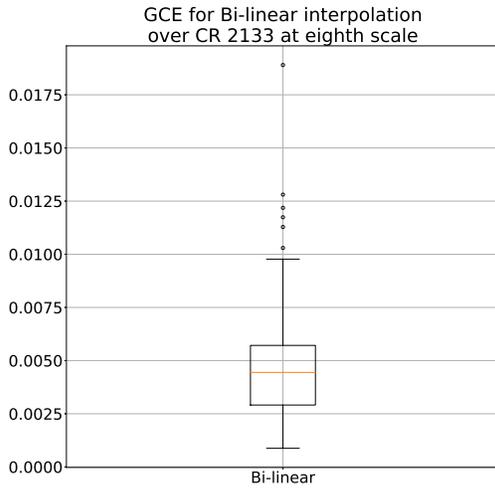
First, regardless of scale, multiple regions within the initial seed often contribute to a single CH or filament detection. The effects of this are partially demonstrated in Figure 2, where the use of HMI magnetogram data minimized evolution of filament regions, resulting in small areas around the initial seeds rather than the unified region seen in the EUV-only segmentations of Grajeda et al. [2023]. Further demonstration of this behavior, however, can be shown by counting the number of regions present in the initial seed and comparing that to a count of the number of regions present in the final EUV-only segmentation. This comparison for CR 2133 is provided in Figure 9, for the full-scale EUV-only segmentations, and Figure 10, for the one-eighth-scale EUV-only segmentations. In both figures the dashed-blue line represents an ideal case where for every region in the initial seed, one region (CH or filament) is produced in the final



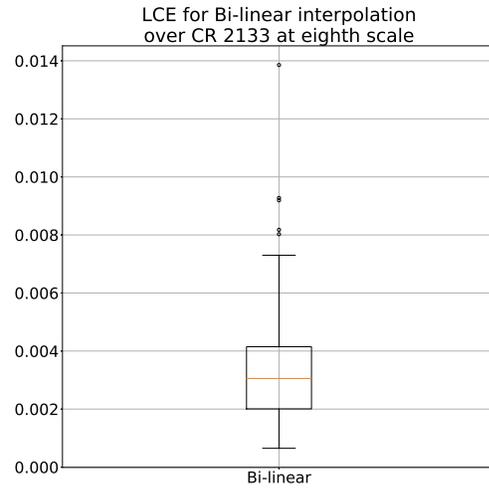
(a) Intersection Over Union



(b) Structural Similarity



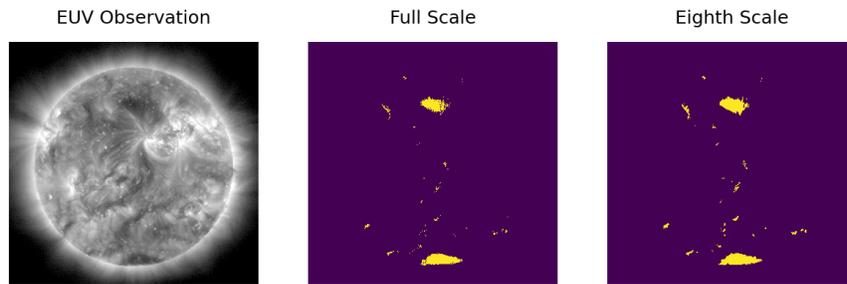
(c) Global Consistency Error



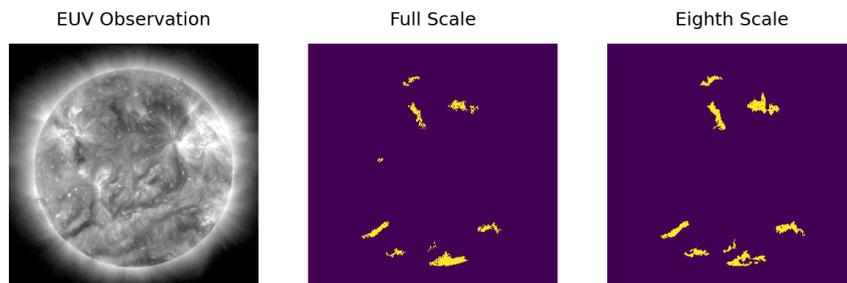
(d) Local Consistency Error

Figure 7: Comparison between full-scale (4096×4096 pixels) segmentation results using both 193 \AA and HMI magnetogram data, and one-eighth-scale (512×512 pixels) segmentation results using both 193 \AA and HMI magnetogram data for CR 2133 when the full-scale seed is utilized to generate the one-eighth-scale segmentation. The box outlines the range between Q1 and Q3, with the median value in orange. The whiskers show 1.5 times the interquartile range. Outliers are marked with circles.

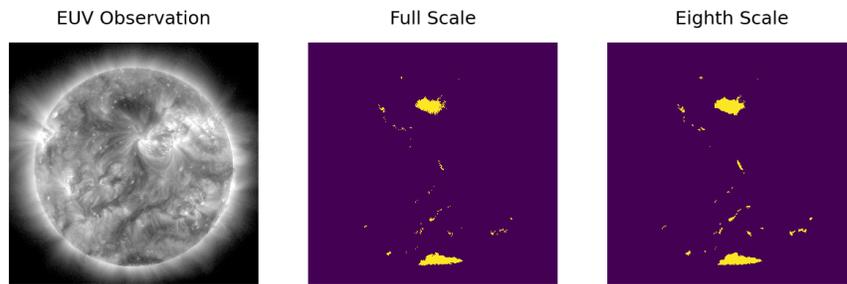
2013-02-12T01:00:01.000



2013-02-13T12:00:01.000



2013-02-11T19:00:01.000



2013-02-12T05:00:01.000

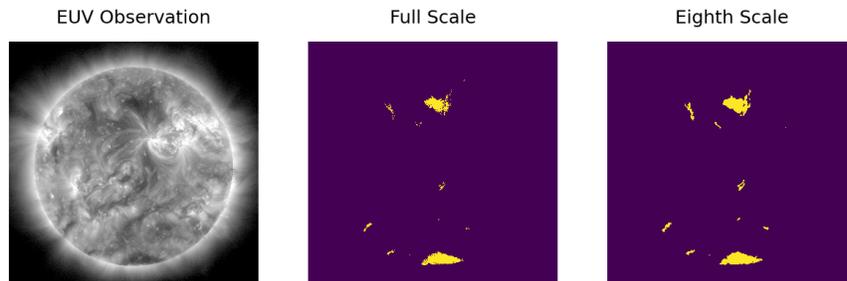


Figure 8: Effects of seed transfer on EUV+HMI magnetogram segmentations generated at one-eighth-scale (512×512 pixels) resolution (right), compared to the corresponding full-scale segmentations (center). In each figure the leftmost image is the 193 \AA observations. The title of each figure is the record time for the EUV data.

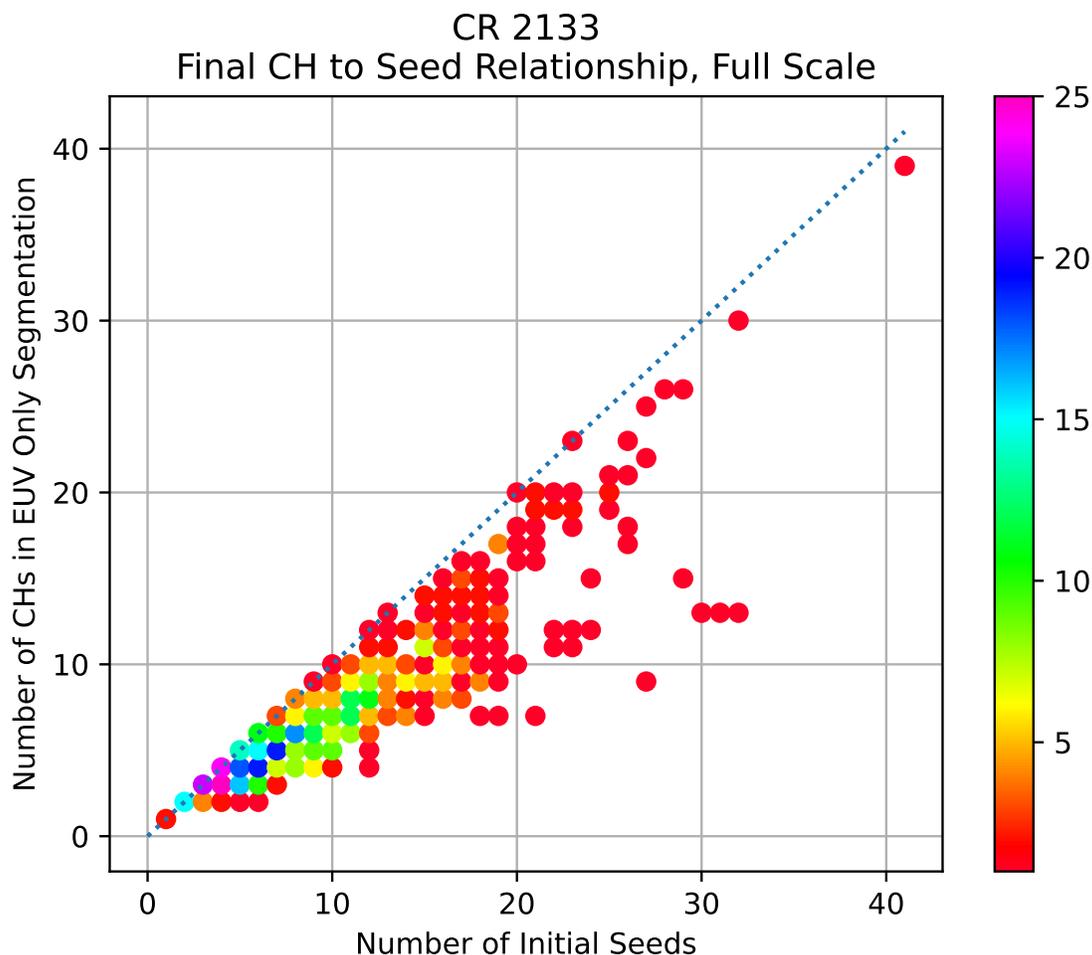


Figure 9: Scatter plot showing the number of regions in the initial seed (x -axis) and number of regions in the final segmentation (y -axis) for EUV-only segmentation at full scale (4096×4096 pixels) for CR 2133. The color of each point denotes the number of images in the dataset that had that specific number of input regions versus that specific number of output regions. The diagonal line is the $x = y$ line, representing the case where for every region in the initial seed there is exactly one region in the final segmentation.

segmentation. Any cases above this line would represent a case where there were more regions in the final segmentation (region splitting), while any case below this represents an example of more regions in the initial seed compared to the final contour (region merging). We note that region splitting does not occur, however region merging is prevalent within the data.

Multiple seeds contributing to a single CH is only problematic due to the fact that it exacerbates the second issue: small-scale statistics. It is reasonable to assume that regions in the initial contour with close physical proximity are part of the same final region, however, we note that regions with a large physical distance between them can belong to the same solar feature. In the counts for Figure 9 and Figure 10, we account for proximity by defining a region in the same manner as we did in Grajeda et al. [2023], namely a region, whether in the initial seed or in the final segmentation, consists of all pixels that will become a single eight-connected component when dilated with either a 40×40 pixel square at full scale, or a 5×5 pixel square at one-eighth scale. This same definition is applied for the plots in Figure 11 (full scale) and Figure 12 (one-eighth scale) which show scatter plots of initial versus final unipolarity for each region in the EUV-only segmentations.

We note that both Figure 11 and Figure 12 indicate regions with a high final (y -axis) value across all initial (x -axis) unipolarity values. This indicates that bipolar regions may appear to have any unipolarity, including high unipolarity, when evaluating the initial seed. We also note, referring to the cluster along the y -axis of Figure 11, that regions

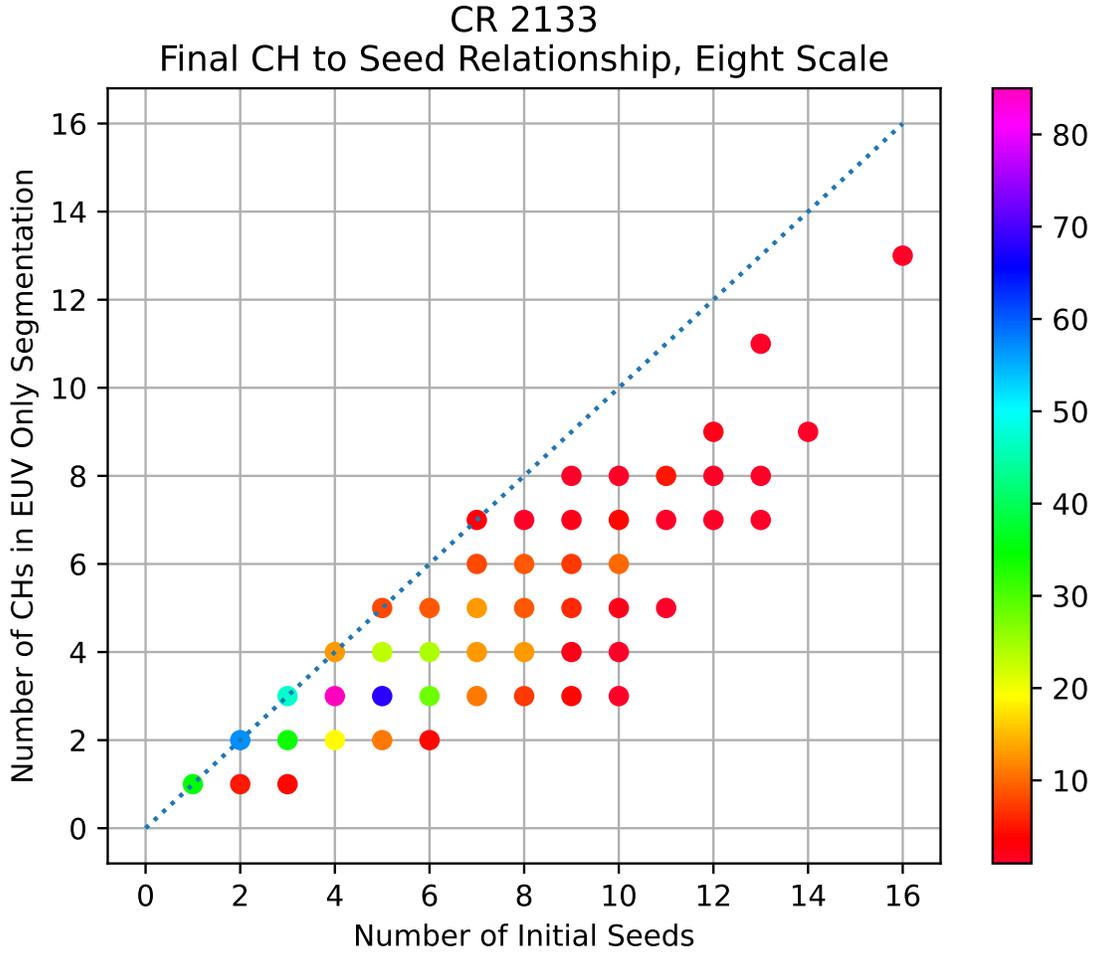


Figure 10: Scatter plot showing the number of regions in the initial seed (x-axis) and number of regions in the final segmentation (y-axis) for EUV-only segmentation at one-eighth scale (512×512 pixels) for CR 2133. The color of each point denotes the number of images in the dataset that had that specific number of input regions versus that specific number of output regions. The diagonal line is the $x = y$ line, representing the case where for every region in the initial seed there is exactly one region in the final segmentation.

reporting high unipolarity at the initial seed are especially prevalent at the full-scale resolution. This explains why the initial seeds in filament regions remain, even when evolving using both EUV and HMI magnetogram observations. The significantly smaller number of observations on the y-axis of Figure 12 provides further indication that seeding at the reduced one-eighth-scale resolution helps to minimize filament contamination.

In summary, we have demonstrated that multiple seeds contribute to a final evolved region (be it a CH or a filament). Furthermore, we have shown that initial seeds may provide an inaccurate estimate of the unipolarity of the final region, mainly due to small-sample statistics. In light of these results, we now consider a means to remove regions associated with filaments from the ACWE evolution.

4.2 Determining Region Evolution

In order to determine if filament regions can be removed during evolution, instead of as a post-hoc process, the unipolarity of every region as a function of number of iterations was calculated for both the EUV-only and EUV + HMI magnetogram segmentations. This was achieved by saving a copy of the contour after each iteration (assuming seeding and iteration at one-eighth-scale resolution). Once the output of each iteration was obtained, the segmentations were

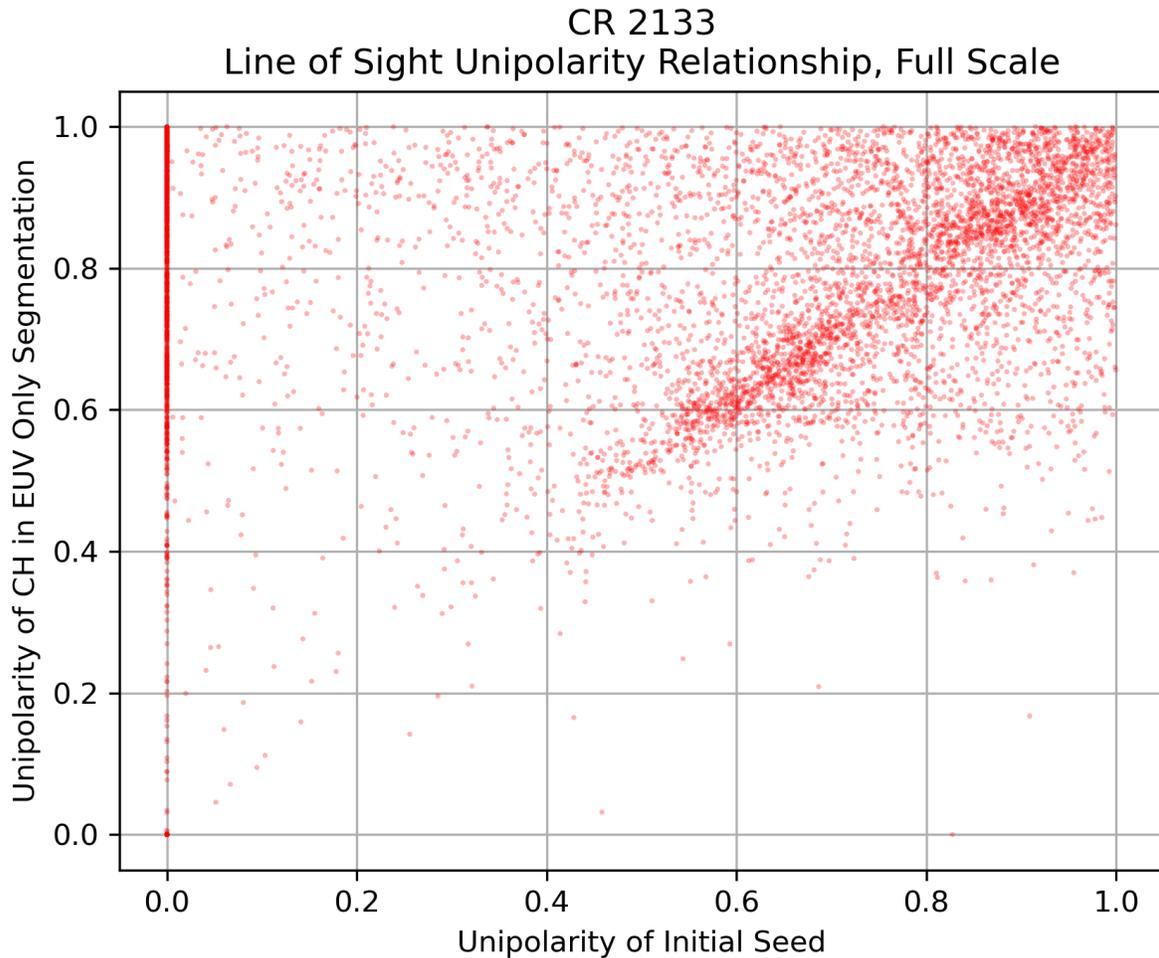


Figure 11: Scatter plot showing the line-of-sight unipolarity of the initial seeds (x-axis) versus the line-of-sight unipolarity of corresponding region in the final segmentation (y-axis) for EUV-only segmentation at full scale (4096×4096 pixels) for all seed regions in CR 2133.

resized to the full-scale resolution, and the line-of-sight unipolarity of each region was calculated using the original, full-scale HMI magnetogram observation.

Figure 13 provides two examples of the change in measured unipolarity of a filament region as a function of ACWE evolution and the corresponding change in area measured in pixels at the full-scale resolution. For both examples we provide the line-of-sight unipolarity, calculated from the HMI magnetogram via Equation 6, and a weighted unipolarity, weighting regions near disk center (where estimates of magnetic field strength are more accurate) using the method outlined in the SunPy documentation [The SunPy Community, 2024] as functions of the number of iterations on the left panel of each example. In the right panel of each example, we provide the area of the region as a function of the number of iterations.

The two examples in Figure 13 were chosen to demonstrate the two primary behaviors noted among filaments in CR 2133. The first case, seen in Figure 13a, consists of regions that quickly become bipolar as they evolve. In these cases, minimal evolution, even with spatially-decimated HMI magnetogram data, causes the region to quickly become bipolar. We note that when this occurs, the spatially-decimated HMI magnetogram further constrains evolution as compared to an EUV-only evolution, preventing further growth of the region. In these cases the filament region can be identified and removed after very few iterations. This case is representative of most filaments found in CR 2133.

The second case, represented in Figure 13b, consists of regions that take several iterations before they can be correctly identified as filament regions. In these cases the spatially-decimated HMI magnetogram data often prevents the region

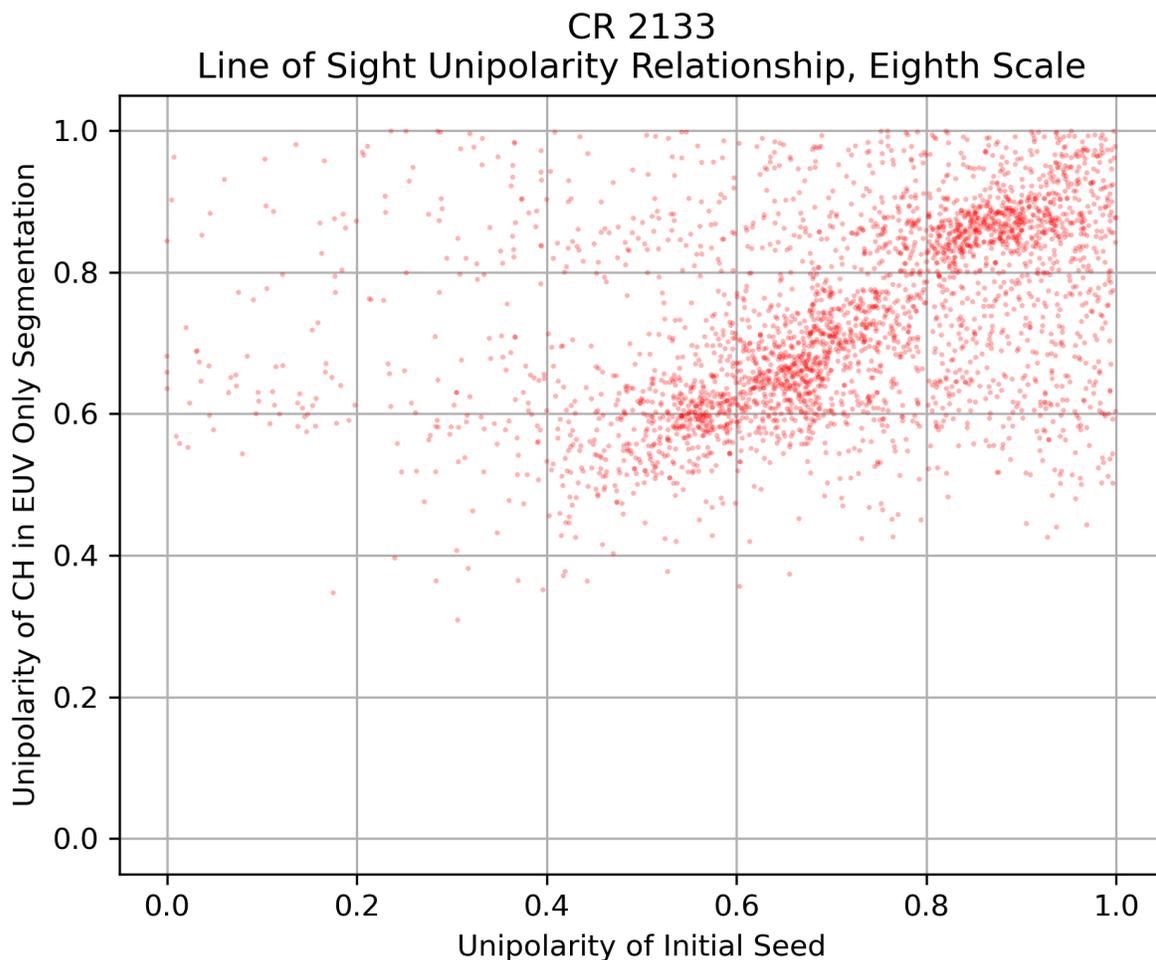


Figure 12: Scatter plot showing the line-of-sight unipolarity of the initial seeds (x-axis) versus the line-of-sight unipolarity of corresponding region in the final segmentation (y-axis) for EUV-only segmentation at one-eighth-scale (512×512 pixels) for all seed regions in CR 2133.

from evolving before it can be correctly identified as a bipolar region. Using EUV only evolution, these cases typically become large enough to correctly identify in < 50 iterations. Due to the presence of filament regions with the behavior seen in Figure 13b, we remove potential filament regions after evolving the contour for 50 EUV-only iterations. To allow for the variation in measured unipolarity seen in both examples in Figure 13 and in other similar cases, we remove regions with a unipolarity $U \geq 0.8$.

Interestingly, we additionally note in Figure 13 that the incorporation of magnetic field appears to result in faster convergence. While definitive reasons for this are unclear, we hypothesize that it is due to a combination of two factors. First, estimates of the CH are now more conservative in general since the evolution is guided away from QS regions by the unipolarity constraint. Second, there are now two forces (homogeneity and unipolarity) that are both acting on the contour, guiding it to a stable (albeit smaller) solution faster.

In summary, inclusion of HMI data and the unipolarity term in the energy functional constrains evolution of filament regions to remain small and isolated around the initial seed (which is determined based solely on EUV intensity). That constraint, however, can result in the area remaining too small for accurate statistics in the computation of unipolarity, resulting in misidentification of filaments as small CHs. These results indicate that a seed filtering process to remove filaments can utilize initial evolution with only EUV data, removal of regions with insufficient unipolarity, and evolution of remaining regions to convergence. This process is elaborated on in the next section.

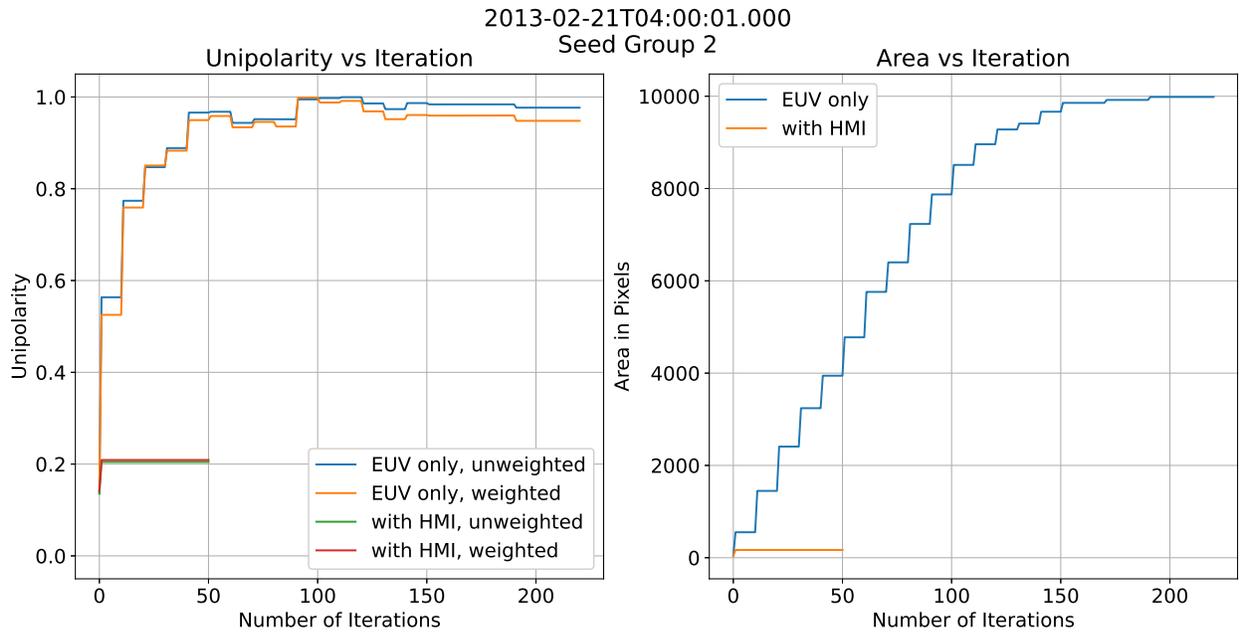
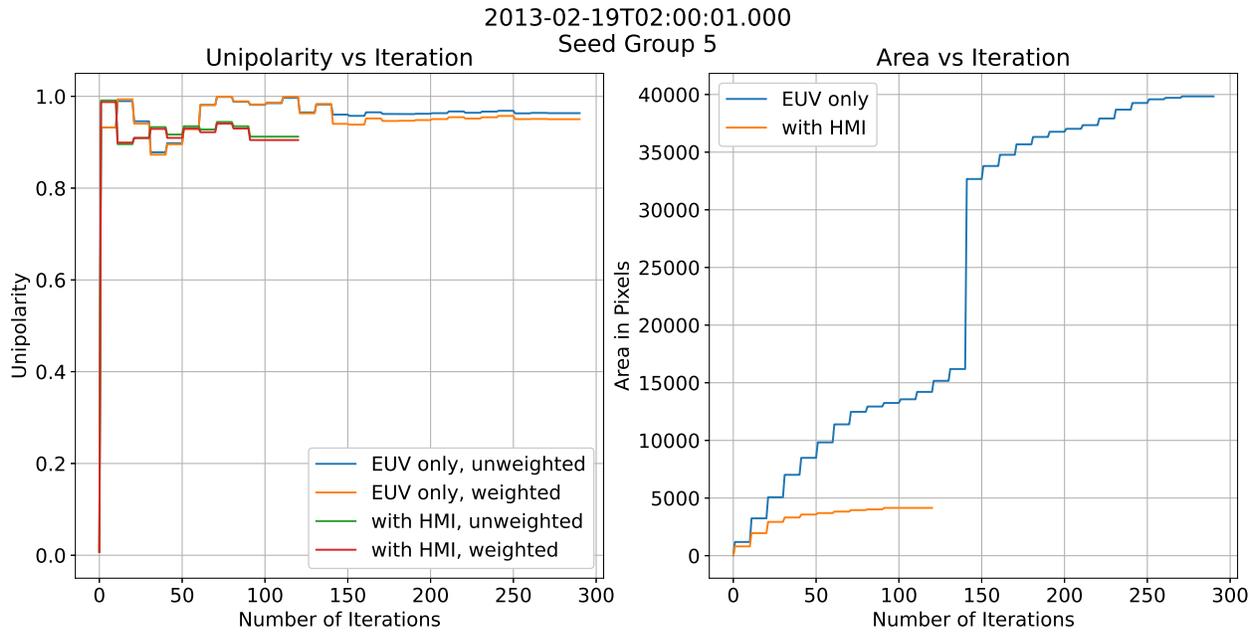


Figure 13: Region unipolarity (left) and region size in pixels (right) as a function of iteration for EUV-only and EUV+HMI magnetogram segmentations generated at one-eighth scale resolution.

5 Updated ACWE Segmentation Method

5.1 Updated Pipeline

Based on the results from Section 4.2, which demonstrate that the initial seed does not provide sufficient information to correctly identify filaments, and the results from Section 3.4, which demonstrate that decimated HMI magnetogram data still constrains CH evolution to highly unipolar regions, we develop the following ACWE pipeline:

1. Open and, if needed, update all observations to ensure correct header information, consistent solar radii, and that all observations are oriented so that solar north is at the top of the image.
2. Resize all observations to 512×512 pixels.
3. Correct EUV observation(s) for limb brightening.
4. Generate initial seed using EUV observation(s).
5. Combine all observations into a vector-valued image.
6. Mask off-disk areas and evolve using only the EUV observation(s) until the unipolarity of each region can be accurately estimated, approximately 50 iterations. These initial iterations set $\lambda_j^{U+} = \lambda_j^{U-} = 0$ where j is the index of the HMI magnetogram data, and evolve using Equation 11 as the energy functional.
7. Identify and remove bipolar regions.
8. Continue evolving using both EUV and HMI magnetogram data (by setting λ_j^{U+} and λ_j^{U-} to non-zero values) until convergence.

We note that every ten iterations, both when evolving only EUV data and when evolving with both EUV and HMI magnetogram data, we reset the intensity of the off-disk regions to the channel-specific mean of the non-CH, on-disk region.

5.2 Results

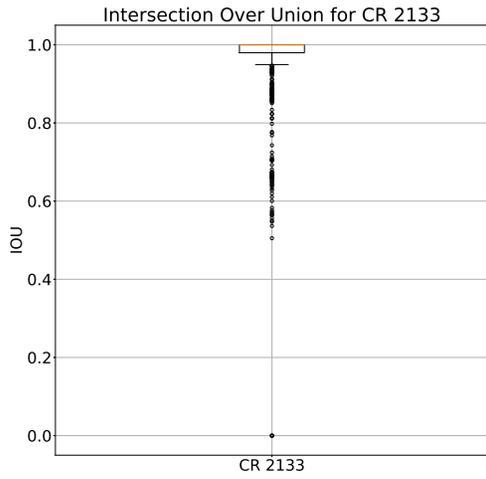
For the results in this section, we use the EUV 193 Å observation, seed using $\alpha = 0.3$, and perform EUV-only evolution for 50 iterations using the homogeneity parameters from Section 3.3 before removing regions with a line-of-sight unipolarity metric ≥ 0.8 . We perform the full EUV+HMI magnetogram evolution on the remaining regions using the weights from Section 3.3 for all parameters. Figure 14 summarizes the similarity between the one-eighth-scale seeded and evolved EUV+HMI magnetogram segmentations in Section 3.4, and the segmentations generated here for CR 2133. We note that the median IOU and SSIM are both one (1) and that the median GCE and LCE are both zero (0), indicating that the vast majority of segmentations are identical. This is consistent with the expectation that seeding at one-eighth-scale already removed the majority of filaments (see Section 3.4). Figure 15 provides samples of the segmentations with differences between the methods. We note that the regions that are missing as a result of filtering the seed are regions that appear highly bipolar based on line-of-sight, full-scale magnetogram data. Referencing Figure 15c, we further note that this can, in rare cases, result in the loss of observations near disk edge. Exploration of a variable unipolarity threshold that is a function of distance from disk center is therefore an area of future research.

Noting that IOU is the most stringent of the four metrics, we provide the comparisons between EUV+HMI magnetogram data with and without seed filtering for all for CRs as box plots in Figure 16 and as 200-bin histograms in Figure 17. Visual inspection of all segmentations with discrepancies revealed that the majority of missing regions consisted of small regions that appear highly bipolar in full-scale line-of-sight magnetogram data. The remaining missing regions consisted of regions along the disk edge. All regions that were retained, in every CR, appeared visually identical.

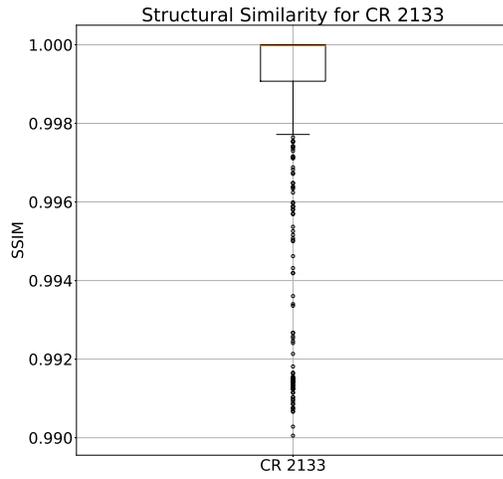
6 Conclusions and Future Work

Through the incorporation of magnetic field information, this new formulation of the ACWE algorithm is able to simultaneously consider region homogeneity in EUV and unipolarity in the magnetic field when delineating between CHs and other solar features. By relying on the fact that CHs appear as dark, homogeneous regions in EUV and as unipolar regions within the magnetic field, this method is able to produce more robust segmentations by retaining CH observations while simultaneously reducing the total area of false positive detections such as surrounding QS (seen in CR 2099) and filaments (seen in CR 2133).

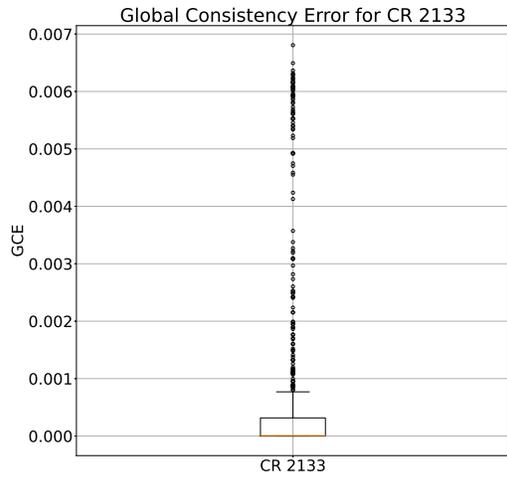
Furthermore, we note that these benefits are retained even with aggressive spatial rescaling of both EUV and HMI magnetogram data to one-eighth of their original size (or 512×512 pixels). This rescaling offers two benefits. First,



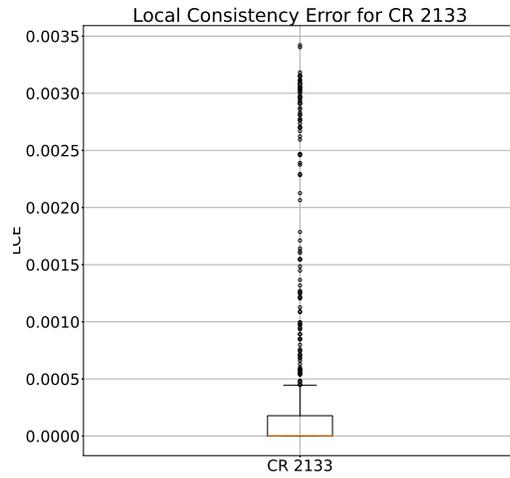
(a) IOU



(b) SSIM



(c) GCE



(d) LCE

Figure 14: Comparison between one-eighth-scale seeded and evolved EUV+HMI magnetogram segmentations and EUV+HMI magnetogram segmentations generated from a filtered seed over CR 2133. The box outlines the range between Q1 and Q3, with the median value in orange. The whiskers show 1.5 times the interquartile range. Outliers are marked with circles.

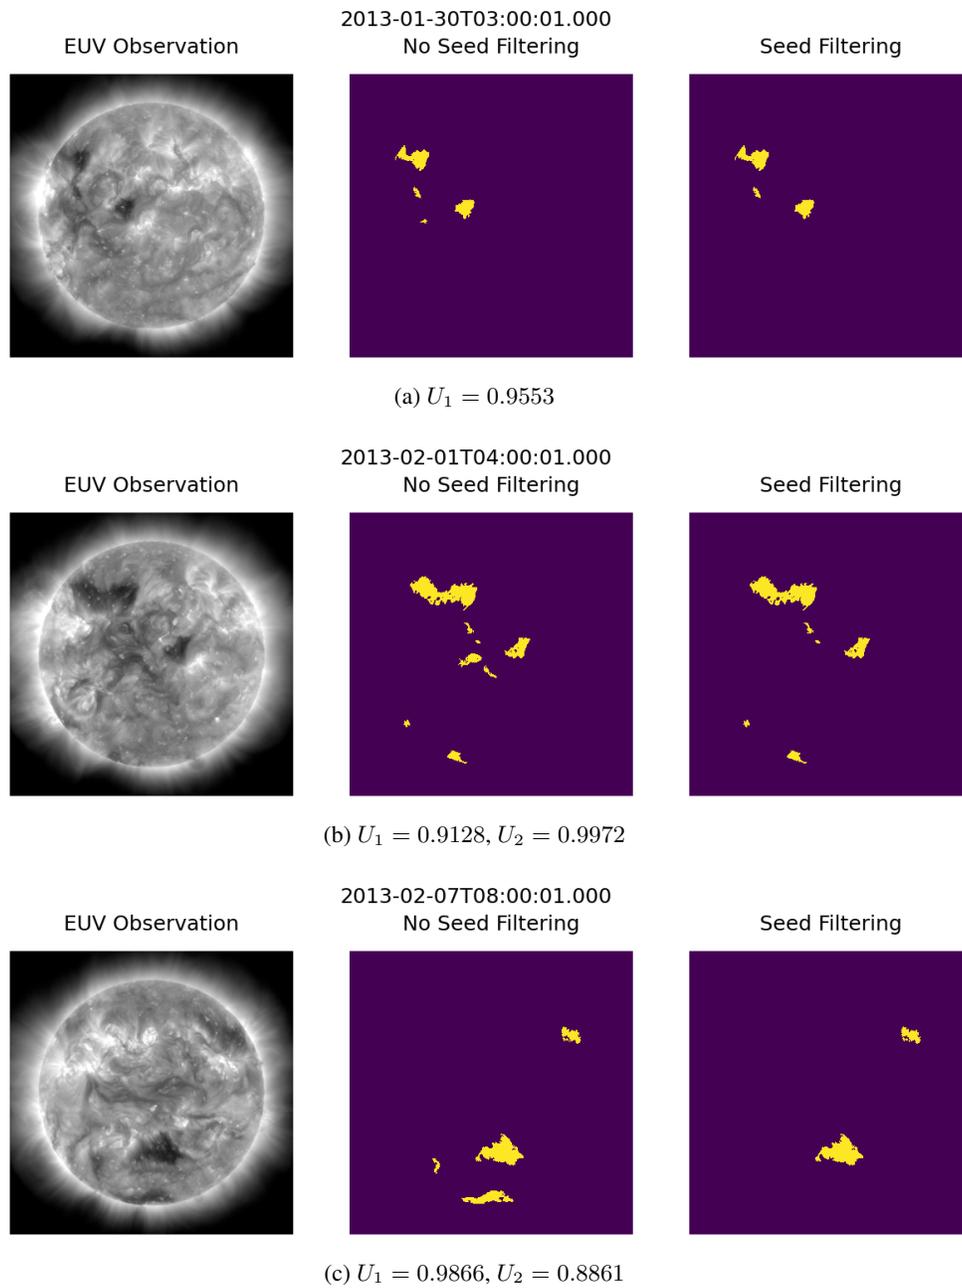
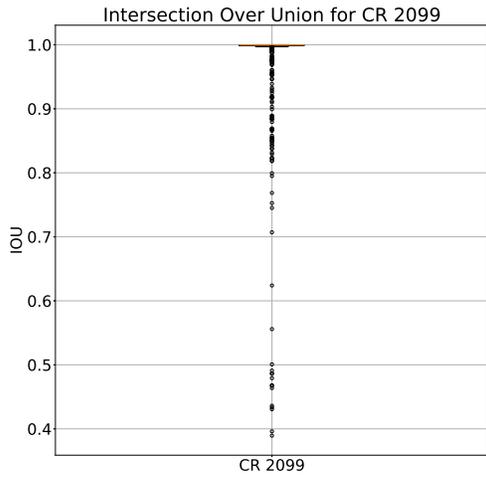
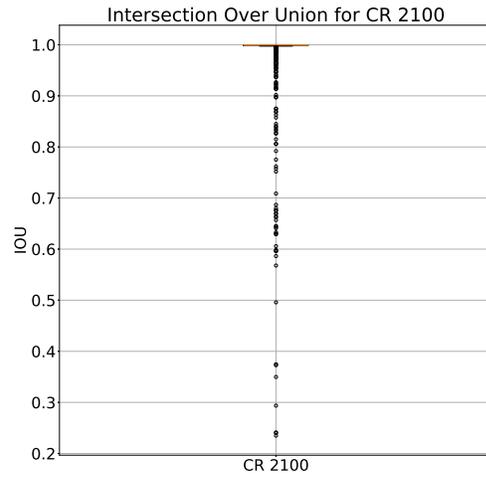


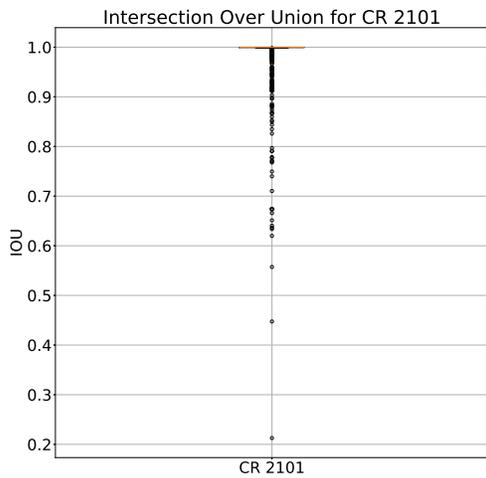
Figure 15: Example of segmentations from CR 2133 generated from EUV and HMI magnetogram observations without filtering the initial seed (center) and with seed filtering (right). The caption under each example provides the unweighted line-of-sight unipolarity of the missing regions in order of right to left, top to bottom.



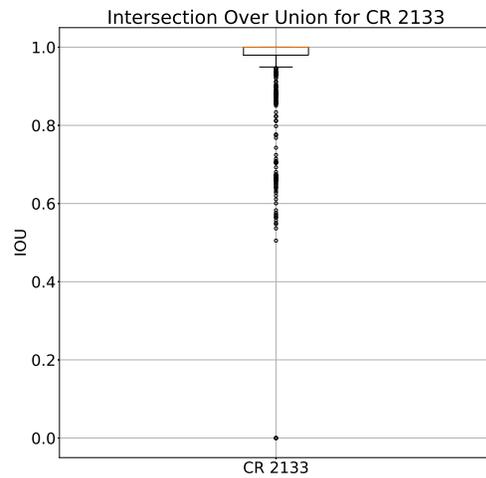
(a) CR 2099



(b) CR 2100

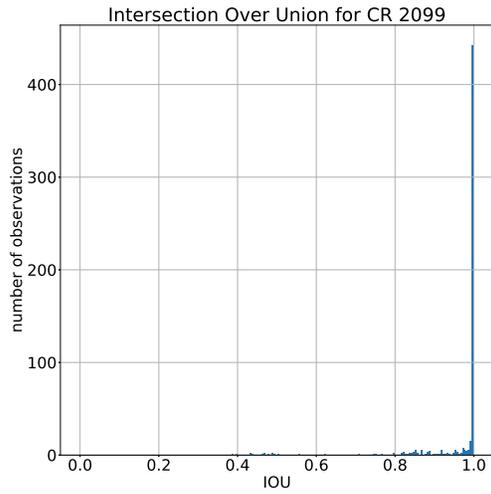


(c) CR 2101

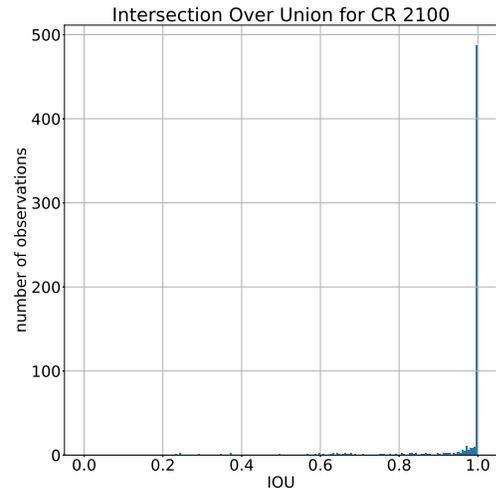


(d) CR 2133

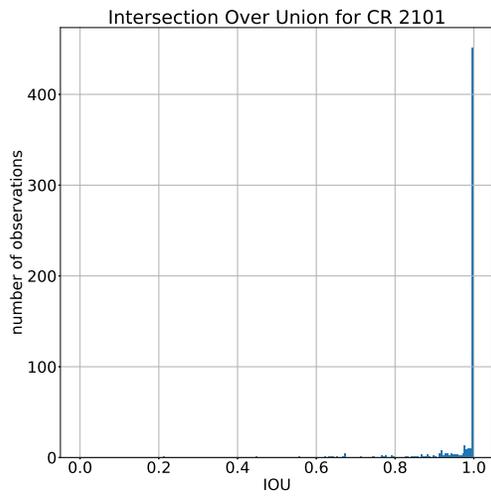
Figure 16: Intersection Over Union between one-eighth-scale seeded and evolved EUV+HMI magnetogram segmentations, and EUV+HMI magnetogram segmentations generated from a filtered seed for each CR. The box outlines the range between Q1 and Q3, with the median value in orange. The whiskers show 1.5 times the interquartile range. Outliers are marked with circles.



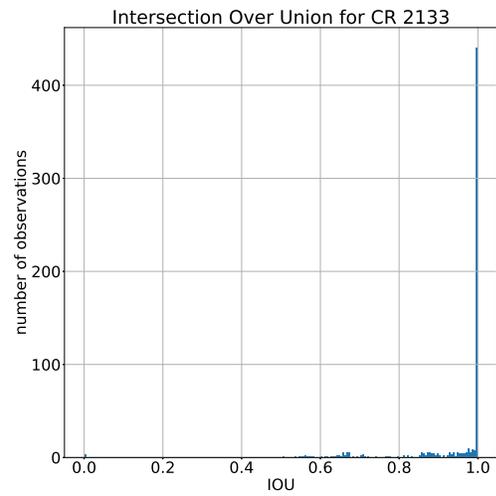
(a) CR 2099



(b) CR 2100



(c) CR 2101



(d) CR 2133

Figure 17: Intersection Over Union between one-eighth-scale seeded and evolved EUV+HMI magnetogram segmentations, and EUV+HMI magnetogram segmentations generated from a filtered seed for each CR.

like with the prior formulation of ACWE for CH detection [Grajeda et al., 2023], the ability to produce meaningful segmentations at a reduced spatial resolution greatly increases algorithm efficiency. Second, we note that operating at a reduced spatial resolution provides an automatic reduction of filament false-positive detections due to the fact that the thin structure of these filaments increases the likelihood that dark pixels will not be retained at the lower resolution, thus eliminating them from the initial threshold-based seed outright. Even when not eliminated, however, the fact that meaningful segmentations are produced at a reduced spatial resolution ensures that filaments that are misidentified in the initial seed are still constrained, thus minimizing overall contamination. With the addition of seed filtering to eliminate non-unipolar regions, filament contamination can be further reduced without compromising the identified CH regions.

With this noted, there are still several areas of potential improvement. First, we note that the seed filtering process herein proposed can eliminate CHs near disk edge that would otherwise be correctly segmented. Addressing this issue may require a variable unipolarity threshold that takes into account the location of the region with respect to disk center. Second, we note that this method relies on line-of-sight magnetic field data, and hope to explore the use of an estimate of radial magnetic field data in future research to determine if this can further improve segmentations near the limb. Finally, we note that the previous formulation of ACWE [Grajeda et al., 2023] provided the ability to quantify the uncertainty of detection based on the homogeneity of any given region compared to adjacent CH regions identified in the initial seeding process. We hope to reintroduce this ability in a future implementation of this expanded ACWE algorithm, leveraging this algorithm’s ability to minimize inclusion of QS regions in order to minimize change of target cases and better characterize the CH region.

Acknowledgements

The authors gratefully acknowledge the support of NASA grant 80NSSC20K0517 which helped support this work.

Data and Code Availability

The data used in this work are publicly available. The GitHub repository at <https://github.com/DuckDuckPig/CH-QUACK> contains code to download the dataset used in this paper, the base ACWE segmentation code, and code to replicate all experiments described herein.

References

- M. D. Altschuler, D. E. Trotter, and F. Q. Orrall. Coronal holes. *Solar Physics*, 26(2):354–365, October 1972. doi: 10.1007/BF00165276.
- E. Antonucci, M. A. Dodero, S. Giordano, V. Krishnakumar, and G. Noci. Spectroscopic measurement of the plasma electron density and outflow velocity in a polar coronal hole. *Astronomy & Astrophysics*, 416:749–758, March 2004. doi: 10.1051/0004-6361:20031650.
- C. N. Arge, D. Odstrcil, V. J. Pizzo, and L. R. Mayer. Improved method for specifying solar wind speed near the Sun. In Marco Velli, Roberto Bruno, Francesco Malara, and B. Bucci, editors, *Solar Wind Ten*, volume 679 of *American Institute of Physics Conference Series*, pages 190–193, September 2003. doi: 10.1063/1.1618574.
- C. N. Arge, J. G. Luhmann, D. Odstrcil, C. J. Schrijver, and Y. Li. Stream structure and coronal sources of the solar wind during the May 12th, 1997 CME. *Journal of Atmospheric and Solar-Terrestrial Physics*, 66(15-16):1295–1309, October 2004. doi: 10.1016/j.jastp.2004.03.018.
- W. T. Barnes, M. C. M. Cheung, M. G. Bobra, P. F. Boerner, G. Chintzoglou, D. Leonard, S. J. Mumford, N. Padmanabhan, A. Y. Shih, N. Shirman, D. Stansby, and P. J. Wright. aiapy: A python package for analyzing solar EUV image data from AIA. *Journal of Open Source Software*, 5(55):2801, 2020. doi: 10.21105/joss.02801.
- L. E. Boucheron, M. Valluri, and R. T. J. McAteer. Segmentation of coronal holes using active contours without edges. *Solar Physics*, 291(8):2353–2372, 2016. doi: <https://doi.org/10.1007/s11207-016-0985-z>.
- R. M. Caplan, C. Downs, and J. A. Linker. Synchronic coronal hole mapping using multi-instrument EUV images: Data preparation and detection method. *The Astrophysical Journal*, 823(1):53, May 2016. ISSN 1538-4357. doi: 10.3847/0004-637x/823/1/53.
- T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001. doi: 10.1109/83.902291.

- T. F. Chan, B. Y. Sandberg, and L. A. Vese. Active contours without edges for vector-valued images. *Journal of Visual Communication and Image Representation*, 11(2):130–141, 2000. ISSN 1047-3203. doi: <https://doi.org/10.1006/jvci.1999.0442>.
- T. M. Garton, P. T. Gallagher, and S. A. Murray. Automated coronal hole identification via multi-thermal intensity segmentation. *Journal of Space Weather and Space Climate*, 8:A02, 2018. doi: 10.1051/swsc/2017039.
- J. A. Grajeda, L. E. Boucheron, M. S. Kirk, A. Leisner, and C. N. Arge. Quantifying the consistency and characterizing the confidence of coronal holes detected by active contours without edges (ACWE). *Solar Physics*, 298(11):133, 2023. doi: <https://doi.org/10.1007/s11207-023-02228-0>.
- A. Hamada, T. Asikainen, I. Virtanen, and K. Mursula. Automated identification of coronal holes from synoptic EUV maps. *Solar Physics*, 293(4):71, April 2018. doi: 10.1007/s11207-018-1289-2.
- R. Jarolim, A. M. Veronig, S. Hofmeister, S. G. Heinemann, M. Temmer, T. Podladchikova, and K. Dissauer. Multi-channel coronal hole detection with convolutional neural networks. *Astronomy & Astrophysics*, 652:A13, 2021. doi: 10.1051/0004-6361/202140640.
- Y.n-K. Ko, K. Muglach, Y.-M. Wang, P. R. Young, and S. T. Lepri. Temporal evolution of solar wind ion composition and their source coronal holes during the declining phase of cycle 23. I. Low-latitude extension of polar coronal holes. *The Astrophysical Journal*, 787(2):121, may 2014. doi: 10.1088/0004-637X/787/2/121.
- L. D. Krista and P. T. Gallagher. Automated coronal hole detection using local intensity thresholding techniques. *Solar Physics*, 256(1–2):87–100, Apr 2009. ISSN 1573-093X. doi: 10.1007/s11207-009-9357-2.
- J. A. Landeros, M. S. Kirk, C. N. Arge, L. E. Boucheron, J. Zhang, V. M. Uritsky, J. A. Grajeda, and M. Dupertuis. Magnetic field-constrained ensemble image segmentation of coronal holes in chromospheric observations, 2024. URL <https://arxiv.org/abs/2405.04731>.
- C. Lowder, J. Qiu, R. Leamon, and Y. Liu. Measurements of EUV coronal holes and open magnetic flux. *The Astrophysical Journal*, 783(2):142, Feb 2014. ISSN 1538-4357. doi: 10.1088/0004-637x/783/2/142.
- D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423 vol.2, 2001. doi: 10.1109/ICCV.2001.937655.
- D. J. McComas, M. Velli, W. S. Lewis, L. W. Acton, M. Balat-Pichelin, V. Bothmer, R. B. Dirling, W. C. Feldman, G. Gloeckler, S. R. Habbal, D. M. Hassler, I. Mann, W. H. Matthaeus, R. L. McNutt, R. A. Mewaldt, N. Murphy, L. Ofman, E. C. Sittler, C. W. Smith, and T. H. Zurbuchen. Understanding coronal heating and solar wind acceleration: Case for in situ near-Sun measurements. *Reviews of Geophysics*, 45(1):RG1004, March 2007. doi: 10.1029/2006RG000195.
- R. H. Munro and G. L. Withbroe. Properties of a coronal “hole” derived from extreme-ultraviolet observations. *The Astrophysical Journal*, 176:511, September 1972. doi: 10.1086/151653.
- M. A. Reiss, K. Muglach, S. Chakraborty, E. Mason, L. Krista, E. Verwichte, C. Foullon, C. Downs, J. Turtle, J. A. Grajeda, L. E. Boucheron, A. Leisner, C. N. Arge, M. S. Kirk, S. Heinemann, R. Jarolim, A. Veronig, V. Delouille, S. Hofmeister, T. M. Garton, C. Lowder, E. Illarionov, Kosovichev A., A. Hamada, C. Möstl, and E. Davies. A community dataset for comparing automated coronal hole detection schemes. *The Astrophysical Journal*, 2023. doi: 10.3847/1538-4365/ad1408.
- T. Robitaille, C. Deil, and A. Ginsburg. reproject: Python-based astronomical image reprojection. <https://ui.adsabs.harvard.edu/abs/2020ascl.soft11023R>, November 2020.
- O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. *Lecture Notes in Computer Science*, page 234–241, 2015. doi: 10.1007/978-3-319-24574-4_28.
- I. F. Scholl and S. R. Habbal. Automatic detection and classification of coronal holes and filaments based on EUV and magnetogram observations of the solar disk. *Solar Physics*, 248(2):425–439, April 2008. doi: 10.1007/s11207-007-9075-6.
- The SunPy Community. Creating a full sun map with AIA and EUVI: Improving the output. https://docs.sunpy.org/en/stable/generated/gallery/map_transformations/reprojection_aia_euvi_mosaic.html#improving-the-output, 2024.
- B. T. Tsurutani, W. D. Gonzalez, A. L. C. Gonzalez, F. L. Guarnieri, N. Gopalswamy, M. Grande, Y. Kamide, Y. Kasahara, G. Lu, I. Mann, et al. Corotating solar wind streams and recurrent geomagnetic activity: A review. *Journal of Geophysical Research: Space Physics*, 111(A7), 2006. doi: <https://doi.org/10.1029/2005JA011273>.
- S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors. scikit-image: image processing in Python. *PeerJ*, 2:e453, 6 2014. ISSN 2167-8359. doi: 10.7717/peerj.453.

- C. Verbeeck, V. Delouille, B. Mampaey, and R. De Visscher. The SPoCA-suite: Software for extraction, characterization, and tracking of active regions and coronal holes on EUV images. *Astronomy & Astrophysics*, 561:A29, 2014. doi: 10.1051/0004-6361/201321243.
- Y. M. Wang and N. R. Sheeley, Jr. Solar wind speed and coronal flux-tube expansion. *The Astrophysical Journal*, 355: 726, June 1990. doi: 10.1086/168805.
- Y.-M. Wang, S. H. Hawley, and N. R. Sheeley, Jr. The magnetic nature of coronal holes. *Science*, 271(5248):464–469, January 1996. doi: 10.1126/science.271.5248.464.
- Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. doi: 10.1109/TIP.2003.819861.