



Tracing H α Fibrils through Bayesian Deep Learning

Haodi Jiang^{1,2} , Ju Jing^{1,3,4} , Jiasheng Wang^{1,3,4} , Chang Liu^{1,3,4} , Qin Li^{1,3,4}, Yan Xu^{1,3,4}, Jason T. L. Wang^{1,2} , and Haimin Wang^{1,3,4}

¹ Institute for Space Weather Sciences, New Jersey Institute of Technology, University Heights, Newark, NJ 07102-1982, USA; hj78@njit.edu, wangj@njit.edu, haimin.wang@njit.edu

² Department of Computer Science, New Jersey Institute of Technology, University Heights, Newark, NJ 07102-1982, USA

³ Center for Solar-Terrestrial Research, New Jersey Institute of Technology, University Heights, Newark, NJ 07102-1982, USA

⁴ Big Bear Solar Observatory, New Jersey Institute of Technology, 40386 North Shore Lane, Big Bear City, CA 92314-9672, USA

Received 2020 November 28; revised 2021 July 10; accepted 2021 July 13; published 2021 September 14

Abstract

We present a new deep-learning method, named FibrilNet, for tracing chromospheric fibrils in H α images of solar observations. Our method consists of a data preprocessing component that prepares training data from a threshold-based tool, a deep-learning model implemented as a Bayesian convolutional neural network for probabilistic image segmentation with uncertainty quantification to predict fibrils, and a post-processing component containing a fibril-fitting algorithm to determine fibril orientations. The FibrilNet tool is applied to high-resolution H α images from an active region (AR 12665) collected by the 1.6 m Goode Solar Telescope (GST) equipped with high-order adaptive optics at the Big Bear Solar Observatory (BBSO). We quantitatively assess the FibrilNet tool, comparing its image segmentation algorithm and fibril-fitting algorithm with those employed by the threshold-based tool. Our experimental results and major findings are summarized as follows. First, the image segmentation results (i.e., the detected fibrils) of the two tools are quite similar, demonstrating the good learning capability of FibrilNet. Second, FibrilNet finds more accurate and smoother fibril orientation angles than the threshold-based tool. Third, FibrilNet is faster than the threshold-based tool and the uncertainty maps produced by FibrilNet not only provide a quantitative way to measure the confidence on each detected fibril, but also help identify fibril structures that are not detected by the threshold-based tool but are inferred through machine learning. Finally, we apply FibrilNet to full-disk H α images from other solar observatories and additional high-resolution H α images collected by BBSO/GST, demonstrating the tool's usability in diverse data sets.

Unified Astronomy Thesaurus concepts: [Solar atmosphere \(1477\)](#); [Solar chromosphere \(1479\)](#); [Convolutional neural networks \(1938\)](#)

1. Introduction

Fibrils are thin threadlike absorption features ubiquitously observed in the solar chromosphere. Depending on their location and dynamic behavior, they may have different names, e.g., threads of filaments (Martin 1998; Wang et al. 2000), the superpenumbra of sunspots (Loughhead 1968; Jing et al. 2019), mottles in quiet-Sun rosette structures (Heinzel & Schmieder 1994), etc. Fibrils are often observed with narrowband solar filtergrams in chromospheric spectral lines such as H α , where they are denser than their surroundings (Moeroogen et al. 2017). Physically speaking, fibrils represent the cool gas “frozen” in magnetic field lines and protected by the magnetic fields from diffusing out (Pikel’ner 1971; Langangen et al. 2008; Rouppe van der Voort et al. 2009). For this reason, fibrils have been traditionally assumed to be aligned with the direction of the chromospheric magnetic field (Foukal 1971a, 1971b).

Tracing chromospheric fibrils in H α is an important subject in heliophysics research (Jing et al. 2011; Leenaarts et al. 2015), and has attracted much attention in the heliophysics community. The comparison between fibrils and the potential magnetic field may provide a quick way to examine the nonpotentiality of active regions (ARs) (Jing et al. 2011). The orientation of fibrils could be used as a constraint to improve the nonlinear force-free modeling of coronal fields (Wiegmann et al. 2008; Aschwanden et al. 2016; Fleishman et al. 2019). Tracing fibrils also helps estimate the amount of

energy in acoustic waves (Fossum & Carlsson 2006) and of free magnetic energy in the chromosphere (Aschwanden et al. 2016).

Many fibril-tracing methods have been developed in recent years. Leenaarts et al. (2015) conducted three-dimensional magnetohydrodynamic simulations to investigate the relation between chromospheric fibrils and magnetic field lines. Aschwanden et al. (2016) performed nonpotential field modeling of chromospheric structures and coronal loops with the VCA-NLFFF code. Jafarzadeh et al. (2017) adopted the CRISPEX tool for visual inspection and identification of isolated slender fibrils. Gafeira et al. (2017) used image processing and contrast enhancement techniques to identify these fibrils. Asensio Ramos et al. (2017) employed the rolling Hough transform (RHT) for fibril detection and a Bayesian hierarchical model to analyze the pixels in spectropolarimetric chromospheric images of penumbrae and fibrils. The authors concluded that fibrils are often well aligned with the magnetic azimuth. This RHT technique has also been used by Schad (2017) to analyze fibrils and coronal rain. Jing et al. (2011) developed a threshold-based algorithm to automatically segment chromospheric fibrils from H α observations and extracted direction information along the fibrils with a fibril-fitting algorithm. The authors further quantitatively measured the nonpotentiality of the fibrils by the magnetic shear angle. In contrast to the above methods, our deep-learning-based tool (FibrilNet) presented here can automatically predict fibrils and measure uncertainties in the predicted results simultaneously.

Deep learning is a branch of machine learning where neural networks are designed to learn from large amounts of data (LeCun et al. 2015). It has been used extensively in computer vision and natural language processing, and more recently in astronomy and astrophysics for flare prediction, spectroscopic analysis, and solar image segmentation, among others (Huertas-Company et al. 2018; Leung & Bovy 2018; Kim et al. 2019; Lieu et al. 2019; Liu et al. 2019; Wu & Boada 2019; Jiang et al. 2020). Different from previous solar image segmentation techniques, which focus on predicting a value for each pixel, our FibrilNet employs a probabilistic segmentation model, specifically a Bayesian convolutional network, that predicts a value for each pixel accompanied with reliable uncertainty quantification. Such a model leads to a more informed decision, and improves the quality of prediction.

In general, there are two types of uncertainties in Bayesian modeling: aleatoric uncertainty and epistemic uncertainty (Kendall & Gal 2017). Aleatoric uncertainty, also known as data uncertainty, measures the noise inherent in observations. Epistemic uncertainty, on the other hand, measures the uncertainty in the parameters of a model; this uncertainty is also known as model uncertainty. Quantifying uncertainties with machine learning finds many applications ranging from computer vision (Kendall & Gal 2017), natural language processing (Xiao & Wang 2019), and medical image analysis (Kwon et al. 2020) to geomagnetic storm forecasting (Gruet et al. 2018; Xu et al. 2020). Here we present a new application of uncertainty quantification with machine learning in fibril tracing.

The rest of this paper is organized as follows. Section 2 describes the solar observations and data used in this study. These data are from an AR (AR 12665) collected by the Big Bear Solar Observatory (BBSO). Section 3 presents the details of our FibrilNet method and the algorithms used by the method. FibrilNet employs a Bayesian convolutional network for probabilistic image segmentation with uncertainty quantification to predict fibrils. It then uses a fibril-fitting algorithm with a polynomial regression function of varying degrees to model the predicted fibrils and determine their orientations. Section 4 reports the experimental results, showing the traced fibrils in $H\alpha$ images of the solar observations in AR 12665 collected by BBSO. Furthermore, we apply FibrilNet to other types of observations, demonstrating the tool's usability in diverse data sets. Section 5 presents a discussion and concludes the paper.

2. Observations and Data Preparation

The Goode Solar Telescope (GST) is a 1.6 m clear-aperture, off-axis telescope at BBSO, which is located in Big Bear Lake, California (Cao et al. 2010; Goode et al. 2010; Goode & Cao 2012; Varsik et al. 2014). GST is equipped with a high-order adaptive optics system, AO-308, which provides high-order correction of atmospheric seeing within an isoplanatic patch (about $6''$ at 500 nm in summer), with a gradual roll-off of correction at larger distances (Shumko et al. 2014). Under a stable seeing condition, BBSO/GST observed AR 12665 at ($W27^\circ$, $S4^\circ$) on 2017 July 13, in which the data taken during $\sim 20:16$ – $22:42$ UT were used in the study presented here.

The Visible Imaging Spectrometer (Cao et al. 2010) of GST utilizes a telecentric mount of the Fabry-Pérot etalon. This imaging system was used for observing the $H\alpha$ line. It scanned the target area at ± 0.6 , ± 0.4 , and 0.0 Å (0.08 Å bandpass) from

the $H\alpha$ line center 6563 Å with a $70''$ circular field of view (FOV). At each wavelength step, the 25 frames, out of 60 frames taken in succession, with the best contrast were saved. These frames, with exposure time ranging from 7 to 20 ms and an image scale of $0''.03$ per pixel, were processed by the high-order adaptive optics system and post-facto speckle image reconstruction algorithms (Wöger et al. 2008), which improved the quality of the images by correcting the wave front deformation caused by atmospheric distortion. An $H\alpha$ line scan was performed over the FOV, and the position with the minimum intensity was defined as the $H\alpha$ line center. It should be pointed out that GST narrowband $H\alpha$ data do not contain the full spectral information, which restricts the full characterization of fibrils in three dimensions. Therefore, our study of fibrils is all based on their projected morphology on the observational image plane.

Our data set contained the GST $H\alpha$ observations in AR 12665 from 20:16:32 UT to 22:41:30 UT on 2017 July 13, where the observed region was located at ($W27^\circ$, $S4^\circ$). During this period of time, 241 $H\alpha$ line center images (i.e., those at 0.0 Å from the $H\alpha$ line center 6563 Å with a $70''$ circular FOV) were used as training data since features in these images were abundant. The test set contained five $H\alpha$ images taken from AR 12665 at 20:15:58 UT on the same day (see Figure 1). Thus, there were 241 training $H\alpha$ images and 5 test $H\alpha$ images, where the size of each image was 720×720 pixels. The training and test sets were disjoint, as the training observations and test observations were taken at different time points. Please note that the five test images were chosen in such a way that they were on five different wavelength positions rather than at five different time points on the same wavelength position. The reason why we did not choose test images that were equally distributed over the time series on the same wavelength position was that the features in the images on the same wavelength position did not change much across the images. By contrast, the features in the images on the five different wavelength positions appeared quite different as shown in Figure 1.

3. Methodology

3.1. Overview of FibrilNet

Figure 2 explains how FibrilNet works. Training $H\alpha$ images are preprocessed in steps 1 and 2, and then used to train the Bayesian deep-learning model (step 3). The trained model takes as input a test $H\alpha$ image (step 4) and produces as output a predicted mask accompanied with results for quantifying aleatoric uncertainty and epistemic uncertainty (step 5). In the post-processing phase (step 6), based on the predicted mask, fibrils on the test $H\alpha$ image are detected and highlighted by thin red curves. Furthermore, the orientations of the detected fibrils are determined based on a fibril-fitting algorithm, where the orientations are shown by different colors.

Specifically, in step 1, we apply the threshold-based tool developed by Jing et al. (2011) to each training $H\alpha$ image described in Section 2 to obtain a corresponding fibril mask. Fibril patterns on this mask are very thick, containing a lot of noise. In step 2, we refine the fibril mask via a skeletonization procedure to obtain a fibril skeleton in which fibrils are marked by black and regions without fibrils are marked by white. The skeletonization procedure works by extracting a region-based shape feature representing the general form of the fibrils. This




Figure 1. Five test images at (A) $0.0 \text{ \AA}$, (B) $+0.4 \text{ \AA}$, (C) $+0.6 \text{ \AA}$, (D) $-0.4 \text{ \AA}$, and (E) $-0.6 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT. Enormous amounts of fibrils exist in these $H\alpha$ images.




Figure 2. Illustration of the proposed method (FibrilNet) for fibril tracing. FibrilNet employs a Bayesian deep-learning model for probabilistic image segmentation with uncertainty quantification to predict fibrils and a fibril-fitting algorithm to determine fibril orientations. The training data used to train the Bayesian deep-learning model are highlighted in the dashed box. The tracing results for the test $H\alpha$ image include the predicted/detected fibrils, their orientations, the aleatoric uncertainty, and the epistemic uncertainty.




Figure 3. Architecture of the encoder–decoder convolutional neural network (i.e., the Bayesian deep-learning model) used in FibrilNet. This network is similar to the one presented in Jiang et al. (2020). See text for their differences.

skeletonization procedure results in better and cleaner images suitable for model training (Umbaugh 2010).

The training $H\alpha$ images and fibril skeletons are then used to train the Bayesian deep-learning model for probabilistic image segmentation and uncertainty quantification (step 3). During training, in order to obtain a robust model, we use the data augmentation technique described in Jiang et al. (2020) to expand the training set by shifting, rotating, flipping, and scaling the training images. In step 4, a test $H\alpha$ image is fed to the trained Bayesian deep-learning model. During testing, we use the Monte Carlo (MC) dropout sampling technique described in Section 3.2 to produce a predicted mask of the test $H\alpha$ image accompanied with aleatoric uncertainty and epistemic uncertainty results (step 5). In step 6, by using the fibril-fitting algorithm based on the polynomial regression model described in Section 3.3, our FibrilNet tool outputs detected fibrils marked by red color on the test $H\alpha$ image and their orientations represented by different colors.

3.2. Implementation of the Bayesian Deep-learning Model in FibrilNet

The Bayesian deep-learning model used by FibrilNet is similar to the model used in SolarUnet (Jiang et al. 2020) for tracking magnetic flux elements. Both models have four encoder blocks (E1, E2, E3, and E4) and four decoder blocks (D1, D2, D3, and D4), mediated by a bottleneck (Bot). See Figure 3 and Jiang et al. (2020) for the configuration and parameter settings of the models. While both models are based on an encoder–decoder convolutional neural network, they differ in three ways. First, in performing 2×2 max pooling, represented by a red arrow in Figure 3, the corresponding max pooling indices are stored. During decoding, the max pooling indices at the corresponding encoder layer are recalled, represented by an orange arrow, to upsample, represented by a green arrow, as done in Badrinarayanan et al. (2017). This upsampling technique used by FibrilNet, designed to reduce the number of trainable parameters in the model (network) and hence save memory, is different from the upconvolution layers used in SolarUnet. Second, since fibril patterns are relatively vague and harder to identify than magnetic flux elements, FibrilNet uses twice as many kernels as SolarUnet in all of the blocks in the encoder and decoder, as well as in the bottleneck. Finally, during testing, instead of using the trained model (network) directly to produce segmentation results as done in

SolarUnet, FibrilNet employs an MC dropout sampling technique, detailed below, to produce, for a test $H\alpha$ image, a predicted mask accompanied with aleatoric uncertainty and epistemic uncertainty results (see Figure 2). This MC dropout sampling technique allows FibrilNet to perform probabilistic image segmentation with uncertainty quantification, which is lacking in SolarUnet.

Specifically, to quantify uncertainty with the convolutional neural network, we use a prior probability, $P(\mathbf{W})$, over the network’s weights, $\mathbf{W}$. During training, pairs of $H\alpha$ images and their corresponding fibril skeletons, collectively referred to as $\mathbf{D}$, are used to train the network. According to Bayes’ theorem,

$$P(\mathbf{W}|\mathbf{D}) = \frac{P(\mathbf{D}|\mathbf{W})P(\mathbf{W})}{P(\mathbf{D})}. \quad (1)$$

Computation of the exact posterior probability, $P(\mathbf{W}|\mathbf{D})$, is intractable (Denker & LeCun 1990). Nevertheless, we can use variational inference (Graves 2011) to learn the variational distribution over the network’s weights parameterized by θ , $q_\theta(\mathbf{W})$, by minimizing the Kullback–Leibler (KL) divergence of $q_\theta(\mathbf{W})$ and $P(\mathbf{W}|\mathbf{D})$ (Blei et al. 2017). It is known that training a network with dropout is equivalent to a variational approximation on the network (Gal & Ghahramani 2016). Furthermore, minimizing the cross-entropy loss of the network is equivalent to the minimization of the KL divergence (Goodfellow et al. 2016). Therefore, we use a binary cross-entropy loss function and the adaptive moment estimation (Adam) optimizer (Goodfellow et al. 2016) with a learning rate of 0.0001 to train our model (network). Let $\hat{\theta}$ denote the optimized variational parameter obtained by training the model (network); we use $q_{\hat{\theta}}(\mathbf{W})$ to represent the optimized weight distribution.

In deep learning, dropout is mainly used to prevent overfitting, where a trained model overfits training data and hence cannot be generalized to make predictions on unseen test data. During training, dropout refers to ignoring or dropping out units (i.e., neurons) of a certain set of neurons that is chosen randomly. During testing, dropout can be used to retrieve T MC samples by processing the input test $H\alpha$ image T times (Gal & Ghahramani 2016). (In the study presented here, T is set to 50.) Each time, a set of weights is randomly drawn from $q_{\hat{\theta}}(\mathbf{W})$. Each pixel in the predicted mask, shown in step 5 of Figure 2, gets a mean and variance over the T samples. If the mean is

greater than or equal to a threshold, the pixel is marked by black indicating that the pixel is part of a fibril; otherwise the pixel is marked by white indicating that the pixel is not part of a fibril. (In the study presented here, the threshold is set to 0.5). Following Kwon et al. (2020), we decompose the variance into aleatoric uncertainty and epistemic uncertainty at the pixel. The aleatoric uncertainty captures the inherent randomness of the predicted result, which comes from the input test H α image, while the epistemic uncertainty comes from the variability of W , which accounts for the uncertainty in the model parameters (weights).

In the post-processing phase, we use a connected-component labeling algorithm (He et al. 2009) to group all adjacent black segments if their pixels in the edges or corners touch each other. For each resulting group, which represents a fibril, we locate its pixels in the predicted mask and highlight their corresponding pixels in the test H α image by red. (Resulting groups containing fewer than 10 pixels are considered noise and filtered out.) We then output the detected fibrils highlighted by red color in the test H α image, as shown in step 6 of Figure 2.

3.3. Implementation of the Fibril-fitting Algorithm in FibrilNet

Most of the detected fibrils are lines or curves. In contrast to Jing et al. (2011), who used a quadratic function to fit the detected fibrils, we adopt a polynomial regression model here. Specifically, our regression model is a polynomial function with varying degrees capable of fitting the detected fibrils with different curvatures. In general, regression analysis investigates the relationship between a dependent variable and an independent variable (Bishop 2006). We model a detected fibril as an n th-degree polynomial function as follows:

$$y = \gamma_0 + \gamma_1 x + \gamma_2 x^2 + \dots + \gamma_n x^n + \epsilon, \quad (2)$$

where γ_i denotes the coefficients and ϵ is a random error term. In Equation (2), the independent variable x represents the x coordinate of a pixel in the detected fibril and the dependent variable y represents the y coordinate of the same pixel, where the x -axis represents the E–W direction and the y -axis represents the S–N direction (see Figure 1). When the degree n equals 1, Equation (2) represents a linear regression model, meaning that the detected fibril is represented by a straight line. In our work, n ranges from 1 to 10.

We then use the least squares method (Ostertagová 2012) to find the optimal γ_i values. There are 10 candidate polynomial functions for representing the detected fibril. We use the R -squared score (Ostertagová 2012) to assess the feasibility of these 10 candidate polynomial functions. Specifically, we choose the candidate polynomial function yielding the largest R -squared score, and use this polynomial function to represent the detected fibril. To determine the orientation of the detected fibril, we calculate the derivative of the chosen polynomial function. For each pixel on the detected fibril, we thus obtain the slope of the tangent at the pixel, leading to the orientation angle of the pixel, denoted by θ_f , with respect to the x -axis. Notice that the orientation angle θ_f is in the 0° – 180° range, as two directions differing by 180° are indistinguishable here because the detected fibril in H α does not carry information on the vertical dimension. Thus, θ_f represents the direction of the detected fibril with a 180° ambiguity (Jing et al. 2011).

4. Results

4.1. Tracing Results of FibrilNet Based on Data from AR 12665

In this series of experiments, we used the 241 H α line center images from 20:16:32 UT to 22:41:30 UT on 2017 July 13 mentioned in Section 2 along with their corresponding fibril skeletons to train the FibrilNet tool as described in Section 3. We then used the trained tool to predict and trace fibrils on the five test images at 0.0 Å, +0.4 Å, +0.6 Å, –0.4 Å, and –0.6 Å from the H α line center 6563 Å with a 70'' circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT (see Figure 1). Figure 4 presents the tracing results on the test image at 0.0 Å; the tracing results on the other four test images can be found in the Appendix. In all of the tracing results, fibrils containing 10 or fewer pixels were treated as noise and excluded.

Figure 4(A) shows the original test H α image. Figure 4(B) shows the enlarged FOV of the region highlighted by white box 1 in Figure 4(A). It can be seen from Figure 4(B) that there are salt-and-pepper noise pixels in the region highlighted by white box 2, which are caused by image reconstruction limitations. Figure 4(C) shows the fibrils (red curves) on the test H α image detected by the tool (after skeletonization) presented in Jing et al. (2011). Figure 4(D) shows the fibrils (red curves) predicted by FibrilNet. FibrilNet uses the images processed by the tool in Jing et al. (2011) as training data. The results in Figures 4(C) and (D) are quite similar, demonstrating the good learning capability of FibrilNet.

Figures 4(E) and (F) show the aleatoric uncertainty (data uncertainty) and epistemic uncertainty (model uncertainty) maps, respectively, produced by FibrilNet. Regions predicted with less uncertainty and higher confidence are colored blue. Regions predicted with more uncertainty and lower confidence are colored red. We can see that the main source of uncertainty is the data rather than the model. Specifically, the values in the data uncertainty map in Figure 4(E) range from 0 to 0.246 while the values in the model uncertainty map in Figure 4(F) range from 0 to 0.086. Furthermore, we observe that the ends of a detected fibril are often associated with higher uncertainty. This happens because there is ambiguity surrounding the transition from the fibril body to the nonfibril background area, a finding consistent with that in object detection with uncertainty quantification reported in the literature (Kendall & Gal 2017; Kwon et al. 2020).

Notice also that the map in Figure 4(E) shows higher uncertainty in the noisy region inside white box 2 compared to that in the region outside white box 2. Specifically, in the noisy region inside white box 2, 90% of the values are contained in the range [0.00014 (5%), 0.20528 (95%)]. By contrast, in the region outside white box 2, 90% of the values are contained in the range [0 (5%), 0.18969 (95%)]. Furthermore, it is observed in Figure 4(C) that the tool in Jing et al. (2011) misses some fibril structures with at least 15 pixels inside white box 3. These fibril structures are not present in the mask predicted by FibrilNet either, as shown inside white box 3 in Figure 4(D). Nevertheless, the uncertainty maps of FibrilNet are able to catch and display these missed fibril structures with higher uncertainty by Bayesian inference, as shown inside white box 3 in Figures 4(E) and (F). This finding demonstrates the usefulness of the uncertainty maps, as they not only provide a quantitative way to measure the confidence on each predicted




Figure 4. Fibril tracing results on the test image at $0.0 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT, where the training data were 241 $H\alpha$ line center images taken from the same AR between 20:16:32 UT and 22:41:30 UT on the same day. (A) The original test $H\alpha$ image. (B) The enlarged FOV of the region highlighted by white box 1 in (A). (C) Fibrils (red curves) on the test $H\alpha$ image detected by the tool in Jing et al. (2011). (D) Fibrils (red curves) on the test $H\alpha$ image predicted by FibrilNet. (E) The aleatoric uncertainty (data uncertainty) map produced by FibrilNet. (F) The epistemic uncertainty (model uncertainty) map produced by FibrilNet.

fibril, but also help identify fibril structures that are not detected by the tool in Jing et al. (2011) but are inferred through machine learning. It should be pointed out that the previous fibril-tracing tool in Jing et al. (2011) does not have the capability of producing these uncertainty maps as described here.

Figure 5 compares the orientation angles of the fibrils found by the tool in Jing et al. (2011) and by FibrilNet. The colors of angles between 0° and 90° range from dark blue to green. The colors of angles between 90° and 180° range from green to dark red. It can be seen from Figure 5 that the orientation angles found by the two tools mostly agree with each other, though the angles detected by FibrilNet tend to be smoother. This happens because FibrilNet uses polynomial functions of varying degrees, as opposed to the quadratic function employed by the tool in Jing et al. (2011), to better fit the detected fibrils with different curvatures. Notice also that the quadratic function used by the tool in Jing et al. (2011) may produce wrong angles, which are calculated correctly by the polynomial regression model of FibrilNet. For example, the orientation angle of the fibril at $E-W = 425''$ and $S-N = -190''$, which is

highlighted by a small red circle in Figure 5, is roughly 90° . It is calculated incorrectly by the tool in Jing et al. (2011), as shown in Figure 5(A). On the other hand, FibrilNet calculates the orientation angle of this fibril correctly, as shown in Figure 5(B). It should be pointed out that the smoother and more accurate orientation angles detected by FibrilNet are due to the better fibril-fitting algorithm used by the tool, as explained above. They are not caused by FibrilNet's Bayesian deep-learning model, whose purpose is mainly image segmentation (i.e., marking each pixel by black indicating the pixel is part of a fibril or by white indicating the pixel is not part of a fibril as shown in the predicted mask in Figure 2) with uncertainty quantification (i.e., producing the uncertainty maps as shown in Figure 4).

4.2. Quantitative Assessment of FibrilNet Based on Data from AR 12665

As mentioned above, FibrilNet has two parts: (i) a Bayesian deep-learning model for predicting fibrils with probabilistic image segmentation, and (ii) a fibril-fitting algorithm for




Figure 5. Orientation angles (colored curves) of the detected fibrils on the test image at $0.0 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT. (A) Fibril orientation angles calculated by the tool in Jing et al. (2011). (B) Fibril orientation angles determined by FibrilNet. The orientation angles of a number of fibrils, some of which are highlighted by small red circles here, are calculated wrongly by the tool in Jing et al. (2011), but correctly by FibrilNet.

determining the orientations of predicted fibrils based on the polynomial regression function in Equation (2). Here, we adopt four measures, defined below, to quantitatively assess the first part, comparing the image segmentation algorithms employed by FibrilNet and the tool (after skeletonization) in Jing et al. (2011), based on the same data from AR 12665 used in Section 4.1. Unlike FibrilNet, which employs deep learning for image segmentation, the tool in Jing et al. (2011) used a threshold-based algorithm rather than machine learning for image segmentation.

Let A (B) denote the set of $720 \times 720 = 518,400$ pixels in the mask (skeleton) predicted by FibrilNet (calculated by the tool in Jing et al. 2011) for a test image. Let $p \in A$ be a pixel in A and let $q \in B$ be p 's corresponding pixel in B , i.e., q is at the same position as p . We use $A \cap_A B$ to represent a subset of pixels in A such that for each pixel p in $A \cap_A B$ and p 's corresponding pixel q in B , p and q are marked by the same color. That is, p and q are both marked by black indicating p and q are part of a fibril in A and B , respectively, or p and q are both marked by white indicating p and q are not part of a fibril in A and B , respectively. Similarly, we use $A \cap_B B$ to represent a subset of pixels in B such that for each pixel q in $A \cap_B B$ and q 's corresponding pixel p in A , q and p are marked by the same color. The first quantitative measure is the pixel similarity (PS), also known as the global accuracy (Badrinarayanan et al. 2017), which is defined as

$$PS = \frac{|A \cap_A B| + |A \cap_B B|}{|A| + |B|}, \quad (3)$$

where $|\cdot|$ is the cardinality of the indicated set. PS is used to assess the pixel-level similarity between the mask A predicted by FibrilNet and the skeleton B calculated by the tool in Jing et al. (2011) for the test image. The value of PS ranges from 0 to 1. The larger (i.e., the closer to 1) the PS value, the greater

the pixel-level similarity between the mask A and the skeleton B .

Let A_F (B_F) denote the set of pixels on the fibrils in A (B). Thus, in A , the pixels in A_F are marked by black while the pixels not in A_F are marked by white. Similarly, in B , the pixels in B_F are marked by black while the pixels not in B_F are marked by white. We use $A_F \cap_{A_F} B_F$ to represent a subset of pixels in A_F such that for each black pixel p in $A_F \cap_{A_F} B_F$, p 's corresponding pixel q is also black, i.e., q is in B_F . Similarly, we use $A_F \cap_{B_F} B_F$ to represent a subset of pixels in B_F such that for each black pixel q in $A_F \cap_{B_F} B_F$, q 's corresponding pixel p is also black, i.e., p is in A_F . The second quantitative measure is the fraction of common fibril pixels (FCFP), defined as

$$FCFP = \frac{|A_F \cap_{A_F} B_F| + |A_F \cap_{B_F} B_F|}{|A_F| + |B_F|}. \quad (4)$$

FCFP is used to measure the pixel-level similarity between the fibrils predicted by FibrilNet and those found by the tool in Jing et al. (2011). The value of FCFP ranges from 0 to 1. The larger (i.e., the closer to 1) the FCFP value, the greater the pixel-level similarity between the fibrils predicted by FibrilNet and those found by the tool in Jing et al. (2011).

The third quantitative measure is the fraction of disjunct fibril pixels (FDFP), defined as

$$FDFP = 1 - FCFP. \quad (5)$$

FDFP is used to measure the pixel-level dissimilarity (distance) between the fibrils predicted by FibrilNet and those found by the tool in Jing et al. (2011). The value of FDFP ranges from 0 to 1. The smaller (i.e., the closer to 0) the FDFP value, the greater the pixel-level similarity between the fibrils predicted by FibrilNet and those found by the tool in Jing et al. (2011).

The fourth quantitative measure is the Rand index (RI; Rand 1971; Unnikrishnan et al. 2005), which calculates the

Table 1

Comparison of the Image Segmentation Algorithms Used in FibrilNet and the Tool of Jing et al. (2011) Based on Four Quantitative Measures and Five Test Images

Test Image	PS	FCFP	FDFP	RI
H α 0.0 Å	0.9576	0.8038	0.1962	0.9188
H α +0.4 Å	0.9571	0.8097	0.1903	0.9178
H α +0.6 Å	0.9659	0.8079	0.1921	0.9340
H α -0.4 Å	0.9546	0.7922	0.2078	0.9134
H α -0.6 Å	0.9536	0.8022	0.1978	0.9115

ratio of pairs of pixels whose colors (black or white) are consistent between the mask A predicted by FibrilNet and the skeleton B calculated by the tool in Jing et al. (2011) for the test image. RI accommodates the inherent ambiguity in image segmentation, and provides region sensitivity and compensation for coloring errors near the ends of detected fibrils. For example, consider a wider fibril. FibrilNet may detect the portion to the left of the center of the fibril and highlight this portion by red. The tool in Jing et al. (2011) may detect the portion to the right of the center of the fibril and highlight that portion by red. Under this circumstance, FCFP does not consider there are common pixels between the two red curves, though RI treats the two red curves as consistent curves. Visually the fibril is indeed found by both tools. As a consequence, RI is often used in comparing image segmentation algorithms. The value of RI also ranges from 0 to 1. The larger (i.e., the closer to 1) the RI value, the greater the visual similarity between the fibrils predicted by FibrilNet and those found by the tool in Jing et al. (2011).

Table 1 presents the quantitative measure values of FibrilNet based on the five test images in Figure 1. It can be seen from the table that the mask predicted by FibrilNet and the skeleton calculated by the tool in Jing et al. (2011) are very similar at pixel level, with PS $\geq 95\%$ on the test images. FCFP is about 80%. However, visually, the similarity/consistency between the fibrils predicted by FibrilNet and those found by the tool in Jing et al. (2011) is much greater, where it is quantitatively assessed with RI $\geq 91\%$ on the test images. This finding is consistent with the results presented in Figures 4(C) and (D).

Next, we quantitatively assess the second part of FibrilNet, comparing the fibril-fitting algorithms employed by FibrilNet and the tool (after skeletonization) in Jing et al. (2011), based on the same data from AR 12665 described in Section 4.1. The fibril-fitting algorithms are used to determine the orientations of detected fibrils. Let θ_f represent the fibril orientation angle of a pixel calculated by the polynomial regression function in FibrilNet, and let θ_j represent the fibril orientation angle of the same pixel calculated by the quadratic function in the tool of Jing et al. (2011). The acute angle difference between θ_f and θ_j , denoted as $\delta(\theta_f, \theta_j)$, is defined as

$$\delta(\theta_f, \theta_j) = \begin{cases} |\theta_f - \theta_j| & \text{if } |\theta_f - \theta_j| \leq 90^\circ \\ 180^\circ - |\theta_f - \theta_j| & \text{otherwise} \end{cases}. \quad (6)$$

The angle difference is decided in favor of an acute or right angle, i.e., $0^\circ \leq \delta(\theta_f, \theta_j) \leq 90^\circ$.

Figure 6 quantitatively compares the orientation angles of common fibril pixels calculated by the fibril-fitting algorithms used in FibrilNet and the tool of Jing et al. (2011) based on the test image at 0.0 Å from the H α line center 6563 Å with a 70'' circular FOV collected in AR 12665 on 2017 July 13 20:15:58

UT. Figure 6(A) shows a 2D histogram of the orientation angles of common fibril pixels produced by the two tools, where the x -axis (y -axis) represents the orientation angles calculated by FibrilNet (the tool of Jing et al. 2011). The 2D histogram is computed by grouping common fibril pixels whose orientation angles are specified by their x and y coordinates into bins, and counting the common fibril pixels in a bin to compute the color of the tile representing the bin. The width of each bin equals 2° . It can be seen from Figure 6(A) that the orientation angles of common fibril pixels calculated by the two tools mostly agree with each other, which is consistent with the findings shown in Figure 5. Figure 6(B) shows the differences of the orientation angles of common fibril pixels produced by the two tools. It can be seen from Figure 6(B) that most of the common fibril pixels have very small orientation angle differences, displayed by a purple color. For the common fibril pixels with large orientation angle differences, the orientation angles calculated by the quadratic function used in the tool of Jing et al. (2011) are often incorrect (see, for example, the fibrils highlighted by small red circles in Figure 6(B) and Figure 5).

4.3. Application of FibrilNet to Other Data

In this series of experiments, we applied FibrilNet to other types of test images, including (i) a full-disk image from the Global Oscillation Network Group (GONG; Harvey et al. 1996; Plowman & Berger 2020) at the National Solar Observatory, (ii) a full-disk image from the Kanzelhöhe Solar Observatory (KSO; Otruba 1999; Otruba et al. 2008), (iii) high-resolution superpenumbral fibrils from BBSO (Jing et al. 2019), and (iv) two high-resolution quiet-Sun regions from BBSO. The GONG full-disk Learmonth reduced H α data in (i) was collected on 2015 September 28 00:01:34 UT. The KSO full-disk raw-image H α data in (ii) was collected on 2015 September 14 09:14:20 UT. The GONG and KSO full-disk images have relatively low resolution. The BBSO superpenumbra of sunspots in (iii) was collected at H α -0.6 Å from AR 12661 (501E, 95N) on 2017 June 4 19:08:44 UT. The two BBSO quiet-Sun regions in (iv) were collected on 2018 July 29 16:33:12 UT and 2020 June 10 16:10:25 UT at H α -0.6 Å from (604E, 125S) and H α 0.0 Å from (283E, 789N), respectively. The FibrilNet tool was trained using the same 241 H α line center images described in Section 2. Here we present results without uncertainty maps. Results with uncertainty maps can be generated as done in Section 4.1.

Figure 7 shows the fibrils (red curves) predicted by FibrilNet on the GONG and KSO test images. Figure 7(A) presents the GONG full-disk H α image. Figure 7(B) shows an enlarged view of the region highlighted by the white box in Figure 7(A). In Figure 7(C), we see that FibrilNet detects many fibrils on the GONG image. Figure 7(D) presents the KSO full-disk H α image. Figure 7(E) shows an enlarged view of the region highlighted by the white box in Figure 7(D). Figure 7(F) clearly demonstrates that FibrilNet detects threads of filaments and fibrils on the KSO image.

Figure 8 presents the fibril prediction results on the BBSO high-resolution test H α images. Figure 8(A) shows the BBSO superpenumbra of sunspots image used in the study. It can be seen that there are superpenumbral fibrils around the sunspot in the center of the image. Figure 8(D) shows the predicted superpenumbral fibrils (red curves) produced by FibrilNet on the image in Figure 8(A). We see in Figure 8(D) that FibrilNet




Figure 6. Quantitative comparison of the orientation angles of common fibril pixels calculated by the fibril-fitting algorithms used in FibrilNet and the tool of Jing et al. (2011) based on the test image at $0.0 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT. (A) 2D histogram of the orientation angles of common fibril pixels produced by the two tools. (B) Differences of the orientation angles of common fibril pixels produced by the two tools. The orientation angles of the fibrils highlighted by small red circles are calculated wrongly by the tool of Jing et al. (2011), but correctly by FibrilNet as indicated in Figure 5.

can distinguish the superpenumbral fibrils from the clusters of spicules nearby. Figures 8(B) and (C) present the two BBSO quiet-Sun regions. Figures 8(E) and (F) show the predicted mottles in the quiet-Sun rosette structures in Figures 8(B) and (C), respectively. These high-resolution $H\alpha$ images clearly demonstrate the good fibril prediction capability of our tool.

5. Discussion and Conclusions

We develop a Bayesian deep-learning method, FibrilNet, for tracing chromospheric fibrils in $H\alpha$ images of solar observations. We apply the FibrilNet tool to high-resolution $H\alpha$ images from an AR (AR 12665) collected by BBSO/GST on 2017 July 13. The tool performs well on these high-resolution $H\alpha$ images, predicting fibrils with uncertainty quantification and determining the orientations of the predicted fibrils. We further apply FibrilNet to full-disk $H\alpha$ images from other solar observatories and additional high-resolution $H\alpha$ images collected by BBSO/GST, demonstrating the tool's usability in diverse data sets.

Our main results are summarized as follows:

1. The encoder-decoder convolutional neural network (i.e., the Bayesian deep-learning model) used in FibrilNet, as illustrated in Figure 3, is an enhancement of two deep-learning models, namely U-Net (Falk et al. 2019), based on which our SolarUnet tool (Jiang et al. 2020) for magnetic tracking was developed, and SegNet (Badrinarayanan et al. 2017). FibrilNet predicts fibrils on a test $H\alpha$ image through image segmentation (i.e., predicting each pixel in the test $H\alpha$ image to be black indicating the pixel is part of a fibril or white indicating the pixel is not part of a fibril). In computer vision and image processing, U-Net and SegNet are two of the best image segmentation models. By combining these two models, FibrilNet produces good image segmentation (i.e., fibril prediction) results, as described in Section 4.
2. The training data set used in this study comprises 241 high-resolution $H\alpha$ line center images in AR 12665 collected by BBSO/GST from 20:16:32 UT to 22:41:30

UT on 2017 July 13. After FibrilNet is trained on this data set, we apply the trained model to predict fibrils on five high-resolution test $H\alpha$ images from the same AR (AR 12665) collected by BBSO/GST on 2017 July 13 20:15:58 UT as described in Section 4.1, as well as on an additional five test $H\alpha$ images including two full-disk $H\alpha$ images from GONG/KSO and three other high-resolution $H\alpha$ images collected by BBSO/GST as described in Section 4.3. Our experimental results show that the Bayesian deep-learning model employed by FibrilNet performs well not only on the five high-resolution test $H\alpha$ images from AR 12665 that are not seen during training, but also on the additional five test $H\alpha$ images. No further training is needed for FibrilNet to predict fibrils in the additional five test $H\alpha$ images. This is achieved by the generalization and inference capabilities of the deep-learning model used by FibrilNet. On the other hand, the threshold-based tool in Jing et al. (2011) is tailored for the high-resolution $H\alpha$ images collected by BBSO/GST. When applying the threshold-based tool in Jing et al. (2011) to the GONG full-disk $H\alpha$ image in Figure 7, the threshold-based tool performs poorly, missing many fibrils on the GONG $H\alpha$ image.

3. FibrilNet obtains training data from the threshold-based tool in Jing et al. (2011), where the training data set contains 241 high-resolution $H\alpha$ line center images from AR 12665 collected by BBSO/GST as described in item 2 above. When applying FibrilNet and the threshold-based tool to the five high-resolution test $H\alpha$ images from the same AR (AR 12665) collected by BBSO/GST, the two tools agree well on the detected fibrils as described in Sections 4.1 and 4.2. This demonstrates the good learning capability of FibrilNet. When predicting fibrils on a test $H\alpha$ image, FibrilNet uses an uncertainty quantification technique (more precisely an MC sampling technique) to process the test $H\alpha$ image T times, where $T = 50$ as described in Section 3.2. Unlike FibrilNet, which employs deep




Figure 7. Fibrils (red curves) predicted by FibrilNet on the GONG and KSO full-disk $H\alpha$ images collected on 2015 September 28 00:01:34 UT and 2015 September 14 09:14:20 UT, respectively. (A) The GONG full-disk $H\alpha$ image. (B) Enlarged view of the region highlighted by the white box in (A). (C) Fibrils predicted by FibrilNet on the image in (B). (D) The KSO full-disk $H\alpha$ image. (E) Enlarged view of the region highlighted by the white box in (D). (F) Fibrils predicted by FibrilNet on the image in (E).

learning, the tool in Jing et al. (2011) used a threshold-based algorithm, rather than machine learning, for image segmentation to detect fibrils on the test $H\alpha$ image. It takes several seconds for the threshold-based tool to process the test $H\alpha$ image. When the uncertainty quantification technique is turned off (i.e., T is set to 1), FibrilNet is 10 times faster than the threshold-based tool in Jing et al. (2011) due to the fact that FibrilNet detects fibrils through making predictions, although the two tools produce similar results. When the uncertainty quantification technique is turned on (i.e., T is set to 50), FibrilNet is as fast as the threshold-based tool while producing uncertainty maps that not only provide a quantitative way to measure the confidence on each detected fibril, but also help identify fibril structures that are not detected by the threshold-based tool (i.e., that do not exist in the training data) but are inferred through machine learning as described in Section 4.1. It is worth noting that the main source of uncertainty is the data rather than our deep-learning model. Uncertainty values are higher in noisy regions of the test $H\alpha$ image. Furthermore, the ends of a predicted fibril are often associated with higher uncertainty, due to the ambiguity surrounding the transition from the fibril body to the nonfibril background area. To the best of our

knowledge, FibrilNet is the first tool capable of predicting fibrils with uncertainty quantification.

4. We conducted additional experiments to evaluate the effectiveness of the data augmentation technique used for training FibrilNet as described in Section 3.1. Our experimental results show that, without the data augmentation technique, the performance of FibrilNet degrades, particularly when the tool is applied to the GONG and KSO full-disk $H\alpha$ images in Figure 7. This happens because the data augmentation technique can increase the generalization and inference capabilities of the Bayesian deep-learning model used by FibrilNet. Our training data set comprises 241 $H\alpha$ line center images from AR 12665 collected on 2017 July 13 as described in Section 2. We also performed experiments where we split the training data set into two parts based on image quality. The first part contained 12 $H\alpha$ line center images with slightly lower quality. The second part contained the remaining 229 $H\alpha$ line center images with higher quality. Since the first part contained too few $H\alpha$ images, we expanded it by including 12 lower-quality images from the other four wavelength positions in AR 12665 studied here, yielding a total of 60 lower-quality $H\alpha$ images. Our experimental results show that the deep-learning models trained by all 241 $H\alpha$ line center images




Figure 8. Fibrils (red curves) predicted by FibrilNet on additional high-resolution BBSO test $H\alpha$ images. (A) The BBSO superpenumbra of sunspots image collected at $H\alpha -0.6 \text{ \AA}$ from AR 12661 (501E, 95N) on 2017 June 4 19:08:44 UT. (B) The BBSO quiet-Sun image collected at $H\alpha -0.6 \text{ \AA}$ from (604E, 125S) on 2018 July 29 16:33:12 UT. (C) The BBSO quiet-Sun image collected at $H\alpha 0.0 \text{ \AA}$ from (283E, 789N) on 2020 June 10 16:10:25 UT. (D) Fibrils predicted by FibrilNet on the image in (A). (E) Fibrils predicted by FibrilNet on the image in (B). (F) Fibrils predicted by FibrilNet on the image in (C).

and by the 229 higher-quality $H\alpha$ line center images produce similar results. On the other hand, the performance of the deep-learning model trained by the 60 lower-quality $H\alpha$ images degrades, and becomes even worse in the absence of data augmentation, particularly when the model is applied to the GONG and KSO full-disk $H\alpha$ images.

5. To further understand the behavior of FibrilNet, we trained the tool using all $241 \times 5 = 1205$ high-resolution $H\alpha$ images from all five wavelength positions in AR 12665 studied here and applied the trained tool to the same test images described in Section 4. The results obtained are similar to those presented here, indicating our tool works equally well with fewer training images. When the tool is trained by a much smaller data set such as one with less than 100 $H\alpha$ line center images from AR 12665 collected on 2017 July 13, the tool still performs well on the high-resolution test $H\alpha$ images described in Section 4, but finds fragmented filaments and fibrils, rather than long, complete filaments and fibrils, on the KSO full-disk $H\alpha$ image in Figure 7, even when the tool is trained by the data augmentation technique with higher-quality training images.
6. As mentioned above, the Bayesian deep-learning model in FibrilNet performs image segmentation to predict fibrils with uncertainty quantification. On the other hand,

the fibril-fitting algorithm in FibrilNet uses a polynomial regression function with varying degrees to calculate the orientation angles of the predicted fibrils. This polynomial regression model produces more accurate and smoother fibril orientation angles than the quadratic function used by the tool in Jing et al. (2011) as described in Sections 4.1 and 4.2. However, if we replace the polynomial regression model by the quadratic function in FibrilNet, the two tools would produce the same orientation angles on common fibril pixels they detect.

We conclude that FibrilNet is an effective and alternative method for fibril tracing. It is expected that this tool will be a useful utility for processing observations from diverse instruments including BBSO/GST and the new Daniel K. Inouye Solar Telescope.

We thank the referee and scientific editor for very helpful and thoughtful comments. We also thank the BBSO/GST team for providing the data used in this study. The BBSO operation is supported by the New Jersey Institute of Technology and U.S. National Science Foundation (NSF) grant AGS-1821294. The GST operation is partly supported by the Korea Astronomy and Space Science Institute, Seoul National

University, and the Key Laboratory of Solar Activities of the Chinese Academy of Sciences (CAS) and the Operation, Maintenance, and Upgrading Fund of CAS for Astronomical Telescopes and Facility Instruments. This work was supported by U.S. NSF grants AGS-1927578 and AGS-1954737. C.L. and H.W. acknowledge the support of NASA under grants 80NSSC20K1282, 80NSSC18K0673, and 80NSSC18K1705.

Facilities: Big Bear Solar Observatory, National Solar Observatory, Kanzelhöhe Solar Observatory.

Appendix

Figures A1–A4 compare the fibril tracing results and fibril orientations obtained by FibrilNet and the tool in Jing et al. (2011)




Figure A1. Fibril tracing results on the test image at $+0.4 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT, where the training data were 241 $H\alpha$ line center images taken from the same AR between 20:16:32 UT and 22:41:30 UT on the same day. (A) Fibrils on the test $H\alpha$ image detected by the tool in Jing et al. (2011). (B) Fibrils on the test $H\alpha$ image predicted by FibrilNet. (C) The aleatoric uncertainty (data uncertainty) map produced by FibrilNet. (D) The epistemic uncertainty (model uncertainty) map produced by FibrilNet. (E) Fibril orientation angles calculated by the tool in Jing et al. (2011). (F) Fibril orientation angles determined by FibrilNet. The orientation angles of a number of fibrils, some of which are highlighted by small red circles here, are calculated wrongly by the tool in Jing et al. (2011), but correctly by FibrilNet.




Figure A2. Fibril tracing results on the test image at $+0.6 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT, where the training data were 241 $H\alpha$ line center images taken from the same AR between 20:16:32 UT and 22:41:30 UT on the same day. (A) Fibrils on the test $H\alpha$ image detected by the tool in Jing et al. (2011). (B) Fibrils on the test $H\alpha$ image predicted by FibrilNet. (C) The aleatoric uncertainty (data uncertainty) map produced by FibrilNet. (D) The epistemic uncertainty (model uncertainty) map produced by FibrilNet. (E) Fibril orientation angles calculated by the tool in Jing et al. (2011). (F) Fibril orientation angles determined by FibrilNet. The orientation angles of a number of fibrils, some of which are highlighted by small red circles here, are calculated wrongly by the tool in Jing et al. (2011), but correctly by FibrilNet.

on the test image at $+0.4 \text{ \AA}$, $+0.6 \text{ \AA}$, $-0.4 \text{ \AA}$, and $-0.6 \text{ \AA}$, respectively, from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT, where

the training data were 241 $H\alpha$ line center images taken from the same AR between 20:16:32 UT and 22:41:30 UT on the same day.




Figure A3. Fibril tracing results on the test image at $-0.4 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT, where the training data were 241 $H\alpha$ line center images taken from the same AR between 20:16:32 UT and 22:41:30 UT on the same day. (A) Fibrils on the test $H\alpha$ image detected by the tool in Jing et al. (2011). (B) Fibrils on the test $H\alpha$ image predicted by FibrilNet. (C) The aleatoric uncertainty (data uncertainty) map produced by FibrilNet. (D) The epistemic uncertainty (model uncertainty) map produced by FibrilNet. (E) Fibril orientation angles calculated by the tool in Jing et al. (2011). (F) Fibril orientation angles determined by FibrilNet. The orientation angles of a number of fibrils, some of which are highlighted by small red circles here, are calculated wrongly by the tool in Jing et al. (2011), but correctly by FibrilNet.




Figure A4. Fibril tracing results on the test image at $-0.6 \text{ \AA}$ from the $H\alpha$ line center $6563 \text{ \AA}$ with a $70''$ circular FOV collected in AR 12665 on 2017 July 13 20:15:58 UT, where the training data were 241 $H\alpha$ line center images taken from the same AR between 20:16:32 UT and 22:41:30 UT on the same day. (A) Fibrils on the test $H\alpha$ image detected by the tool in Jing et al. (2011). (B) Fibrils on the test $H\alpha$ image predicted by FibrilNet. (C) The aleatoric uncertainty (data uncertainty) map produced by FibrilNet. (D) The epistemic uncertainty (model uncertainty) map produced by FibrilNet. (E) Fibril orientation angles calculated by the tool in Jing et al. (2011). (F) Fibril orientation angles determined by FibrilNet. The orientation angles of a number of fibrils, some of which are highlighted by small red circles here, are calculated wrongly by the tool in Jing et al. (2011), but correctly by FibrilNet.

ORCID iDs

Haodi Jiang  <https://orcid.org/0000-0001-6460-408X>
 Ju Jing  <https://orcid.org/0000-0002-8179-3625>
 Jiasheng Wang  <https://orcid.org/0000-0001-5099-8209>
 Chang Liu  <https://orcid.org/0000-0002-6178-7471>
 Jason T. L. Wang  <https://orcid.org/0000-0002-2486-1097>
 Haimin Wang  <https://orcid.org/0000-0002-5233-565X>

References

- Aschwanden, M. J., Reardon, K., & Jess, D. B. 2016, *ApJ*, **826**, 61
- Asensio Ramos, A., de la Cruz Rodríguez, J., Martínez González, M. J., & Socas-Navarro, H. 2017, *A&A*, **599**, A133
- Badrinarayanan, V., Kendall, A., & Cipolla, R. 2017, *ITPAM*, 39, 2481
- Bishop, C. M. 2006, *Pattern Recognition and Machine Learning* (Berlin: Springer)
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. 2017, *J. Am. Stat. Assoc.*, **112**, 859
- Cao, W., Gorceix, N., Coulter, R., et al. 2010, *AN*, **331**, 636
- Denker, J. S., & LeCun, Y. 1990, in *Proc. 3rd Int. Conf. on Neural Information Processing Systems*, NIPS'90, ed. R. Lippmann et al. (San Francisco CA: Morgan Kaufmann), 853
- Falk, T., Mai, D., Bensch, R., et al. 2019, *Nat. Methods*, **16**, 67
- Fleishman, G., Mysh'yakov, I., Stupishin, A., Loukitcheva, M., & Anfinogentov, S. 2019, *ApJ*, **870**, 101
- Fossum, A., & Carlsson, M. 2006, *ApJ*, **646**, 579
- Foukal, P. 1971a, *SoPh*, **20**, 298
- Foukal, P. 1971b, *SoPh*, **19**, 59
- Gafeira, R., Lagg, A., Solanki, S. K., et al. 2017, *ApJS*, **229**, 6
- Gal, Y., & Ghahramani, Z. 2016, in *Proc. 33rd Int. Conf. on Machine Learning*, ed. M.-F. Balcan & K. Q. Weinberger (ML Research Press), 1050
- Goode, P. R., & Cao, W. 2012, *Proc. SPIE*, **8444**, 844403
- Goode, P. R., Yurchyshyn, V., Cao, W., et al. 2010, *ApJL*, **714**, L31
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, *Deep Learning* (Cambridge, MA: MIT Press), <http://www.deeplearningbook.org>
- Graves, A. 2011, in *Advances in Neural Information Processing Systems 24*, ed. J. Shawe-Taylor et al., <http://papers.nips.cc/paper/4329-practical-variational-inference-for-neural-networks>
- Gruet, M. A., Chandorkar, M., Sicard, A., & Camporeale, E. 2018, *SpWea*, **16**, 1882
- Harvey, J. W., Hill, F., Hubbard, R. P., et al. 1996, *Sci*, **272**, 1284
- He, L., Chao, Y., Suzuki, K., & Wu, K. 2009, *PatRe*, **42**, 1977
- Heinzl, P., & Schmieder, B. 1994, *A&A*, **282**, 939
- Huertas-Company, M., Primack, J. R., Dekel, A., et al. 2018, *ApJ*, **858**, 114
- Jafarzadeh, S., Solanki, S. K., Gafeira, R., et al. 2017, *ApJS*, **229**, 9
- Jiang, H., Wang, J., Liu, C., et al. 2020, *ApJS*, **250**, 5
- Jing, J., Li, Q., Liu, C., et al. 2019, *ApJ*, **880**, 143
- Jing, J., Yuan, Y., Reardon, K., et al. 2011, *ApJ*, **739**, 67
- Kendall, A., & Gal, Y. 2017, *Proc. 31st Int. Conf. on Neural Information Processing Systems*, NIPS'17 (Red Hook, NY: Curran Associates Inc.), 5580
- Kim, T., Park, E., Lee, H., et al. 2019, *NatAs*, **3**, 397
- Kwon, Y., Won, J.-H., Kim, B. J., & Paik, M. C. 2020, *Comput. Stat. Data Anal.*, **142**, 106816
- Langangen, Ø., De Pontieu, B., Carlsson, M., et al. 2008, *ApJL*, **679**, L167
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, **521**, 436
- Leenaarts, J., Carlsson, M., & Rouppe van der Voort, L. 2015, *ApJ*, **802**, 136
- Leung, H. W., & Bovy, J. 2018, *MNRAS*, **483**, 3255
- Lieu, M., Conversi, L., Altieri, B., & Carry, B. 2019, *MNRAS*, **485**, 5831
- Liu, H., Liu, C., Wang, J. T. L., & Wang, H. 2019, *ApJ*, **877**, 121
- Loughhead, R. E. 1968, *SoPh*, **5**, 489
- Martin, S. F. 1998, *SoPh*, **182**, 107
- Moeroogen, K., Morton, R. J., & Henriques, V. 2017, *A&A*, **607**, A46
- Ostertagová, E. 2012, *Procedia Eng.*, **48**, 500
- Otruba, W. 1999, in *ASP Conf. Ser. 184, Third Advances in Solar Physics Euroconference: Magnetic Fields and Oscillations*, ed. B. Schmieder, A. Hofmann, & J. Staude (San Francisco, CA: ASP), 314
- Otruba, W., Freislich, H., & Hanslmeier, A. 2008, *CEAB*, **32**, 1
- Pikel'ner, S. B. 1971, *SoPh*, **20**, 286
- Plowman, J. E., & Berger, T. E. 2020, *SoPh*, **295**, 143
- Rand, W. M. 1971, *J. Am. Stat. Assoc.*, **66**, 846
- Rouppe van der Voort, L., Leenaarts, J., de Pontieu, B., Carlsson, M., & Vissers, G. 2009, *ApJ*, **705**, 272
- Schad, T. 2017, *SoPh*, **292**, 132
- Shumko, S., Gorceix, N., Choi, S., et al. 2014, *Proc. SPIE*, **9148**, 914835
- Umbaugh, S. E. 2010, *Digital Image Processing and Analysis: Human and Computer Vision Applications with CVIptools* (2nd ed.; Boca Raton, FL: CRC Press, Inc.)
- Unnikrishnan, R., Pantofaru, C., & Hebert, M. 2005, in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'05)* (Los Alamitos, CA: IEEE Computer Society), 34
- Varsik, J., Plymate, C., Goode, P., et al. 2014, *Proc. SPIE*, **9147**, 91475D
- Wang, J., Li, W., Denker, C., et al. 2000, *ApJ*, **530**, 1071
- Wiegelmann, T., Thalmann, J. K., Schrijver, C. J., De Rosa, M. L., & Metcalf, T. R. 2008, *SoPh*, **247**, 249
- Wöger, F., von der Lühe, O., & Reardon, K. 2008, *A&A*, **488**, 375
- Wu, J. F., & Boada, S. 2019, *MNRAS*, **484**, 4683
- Xiao, Y., & Wang, W. Y. 2019, *Proc. AAAI Conf. Artif. Intell.*, **33**, 7322
- Xu, S. B., Huang, S. Y., Yuan, Z. G., Deng, X. H., & Jiang, K. 2020, *ApJS*, **248**, 14