

# Human-efficient labeling of a solar flux emergence video dataset by a deep learning model

Subhamoy Chatterjee (✉ [subhamoy@boulder.swri.edu](mailto:subhamoy@boulder.swri.edu))

Southwest Research Institute <https://orcid.org/0000-0002-5014-7022>

Andres Munoz-Jaramillo

Southwest Research Institute <https://orcid.org/0000-0002-4716-0840>

Derek Lamb

Southwest Research Institute

---

## Article

**Keywords:** machine learning, convolutional neural networks, labeling datasets

**Posted Date:** May 27th, 2021

**DOI:** <https://doi.org/10.21203/rs.3.rs-366908/v1>

**License:**  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

---

# Human-efficient labeling of a solar flux emergence video dataset by a deep learning model

Subhamoy Chatterjee<sup>1,\*,+</sup>, Andrés Muñoz-Jaramillo<sup>1,+</sup>, and Derek A. Lamb<sup>1</sup>

<sup>1</sup>Southwest Research Institute, Boulder, CO 80302, USA

<sup>+</sup>these authors contributed equally to this work

<sup>\*</sup>[subhamoy@boulder.swri.edu](mailto:subhamoy@boulder.swri.edu)

## ABSTRACT

Machine learning is becoming a critical tool for interrogation of large complex data. However, labeling large datasets is time consuming. Here we show that convolutional neural networks (CNNs), trained on crudely labeled astronomical videos, can be leveraged to improve the quality of data labeling and reduce the need for human intervention. We use videos of the solar photospheric magnetic field, crudely labeled into two classes: emergence or non-emergence of large bipolar magnetic regions (BMRs). We train the CNN using crude labeling, manually verify, correct labeling vs. CNN disagreements, and repeat this process until convergence. This results in a high-quality labeled dataset requiring the manual verification of only  $\sim 50\%$  of all videos. Furthermore, by gradually masking the videos and looking for maximum change in CNN inference, we locate BMR emergence time without retraining the CNN. This demonstrates the versatility of CNNs for simplifying the challenging task of labeling complex dynamic events.

## Introduction

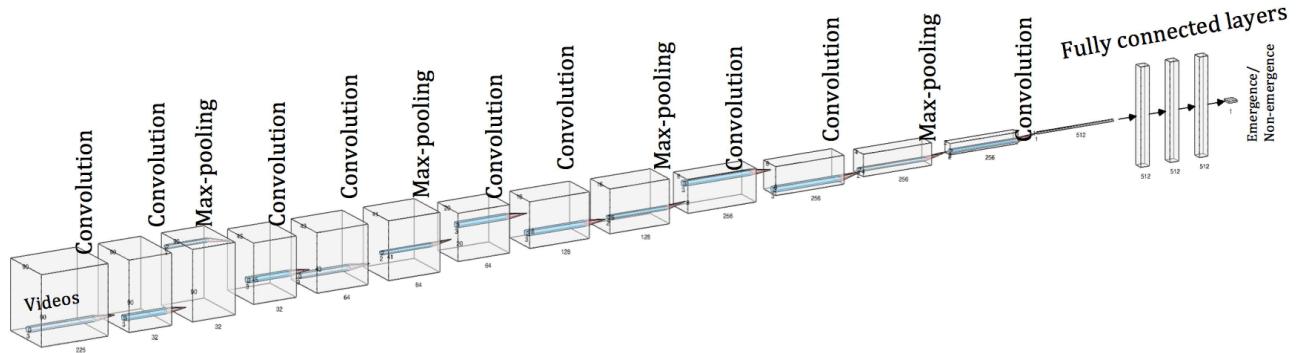
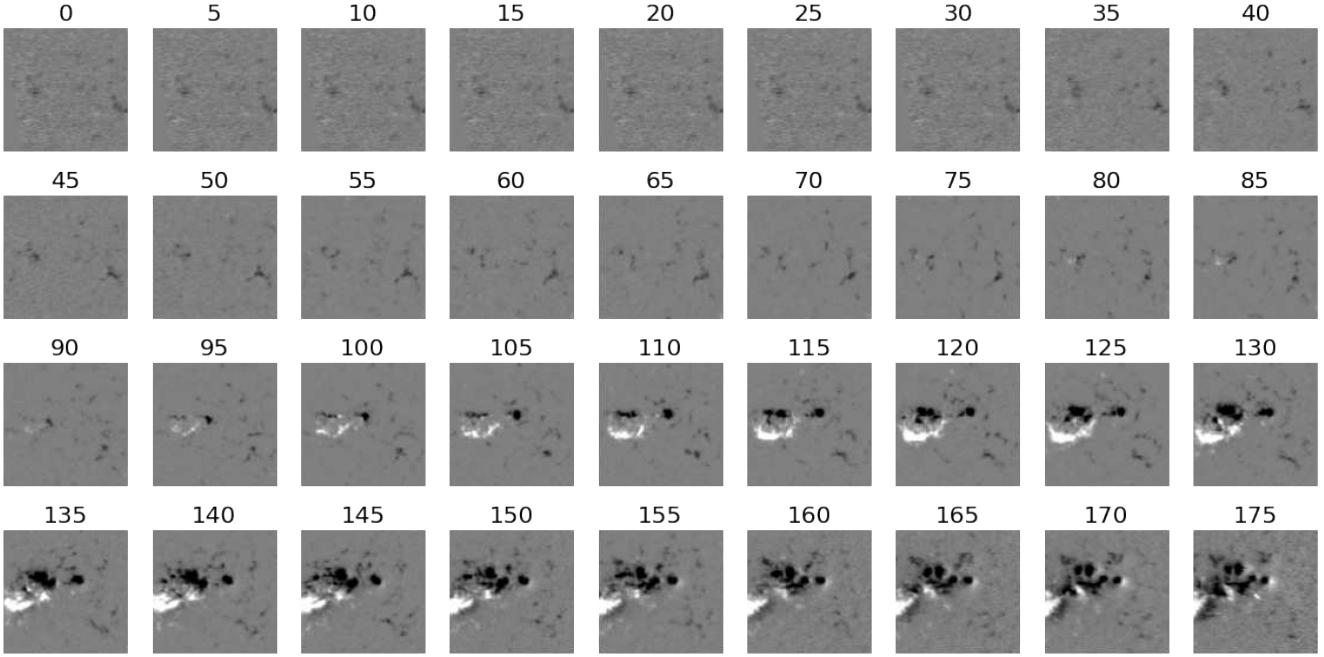
Big-data problems have become increasingly common in astronomy<sup>1</sup>. Large datasets present complex challenges that cannot be tackled with traditional computational techniques. Supervised machine learning (ML), e.g. deep learning, is a promising and effective technique for the classification of complex data such as images and videos<sup>2</sup>. However, manually labelling large databases is a laborious process that requires time and consistency. Iterative labelling approaches, such as ‘active learning’<sup>3</sup> can significantly save time, reducing the cost of making big-data ML ready.

‘Bipolar flux emergence’, which involves the appearance of bipolar/complex magnetic regions (BMRs) on the solar surface, is an example of complex dynamical events that are difficult to label. These regions have the potential to drive space weather events such as flares and coronal mass ejections that can negatively affect satellite networks and long-distance communication<sup>4–6</sup>. Techniques for detecting flux emergence and interactions typically rely on image-by-image processing to perform segmentation and tracking of magnetic elements for detecting dynamic phenomena such as appearance, disappearance, splitting and merging<sup>7–11</sup>. Recently, the single-image segmentation component of this process has been performed using CNNs<sup>12</sup>. However, the complex interaction of magnetic elements demands increasingly complex tracking codes and limits the possibility of real-time detection. For this reason, we have created an end-to-end deep-learning approach for classifying videos of magnetic patch evolution without explicitly supplying segmented images, tracking algorithms or other hand-crafted features. The idea is to allow relevant information abstraction to autonomously take place deeper in the CNN layers to detect flux elements and describe their interaction towards appropriate classification. We start with a crudely labelled solar flux emergence dataset, train a deep learning model, refine the labelling using the trained model and, finally, show how a simple change in data input allows us to detect the time of emergence.

## Results

We use 96-minute cadence videos of SoHO/MDI<sup>13</sup> Line-of-Sight (LoS) magnetic patches ( $15^\circ \times 15^\circ$  in Carrington grid) (Figure 1) labelled for BMR ‘emergence’ (i.e. BMRs that clearly emerge within the visible solar disk) or BMR ‘non-emergence’ (i.e. BMRs that rotate into view). We use BMRs from the Bi-polar Active Region Detection (BARD) code<sup>14</sup> to produce the initial labels. Our initial emergence vs. non-emergence labels are crude in nature, based simply on the first observation of each BMR in the catalog. Any BMR whose first observation occurs to the west of  $-60^\circ$  heliographic longitude, with respect to the solar central meridian, is labeled as an emergence and the rest of as non-emergence. We use the following steps to bring uniformity in the videos as input to a classification model:

1. Missing frames are interpolated for each pixel over the time axis.



**Figure 1.** Dataset and model architecture. *Top panel* shows every 5th frame of an emergence video starting on November 29, 1998 17:36 hours (UT). *Bottom panel* shows adapted deep learning model for classifying videos into ‘emergence’ or ‘non-emergence’ classes. We use repeated convolution, max-pooling and fully connected layers.

- 38     2. Regions outside the solar limb are replaced with a background field of 0 Gauss  
 39     3. Videos are padded with 0 Gauss frames towards the end to make all videos have a standard duration of 225 frames.

40 This translates to a video classification problem with inputs of size  $90 \times 90$  pixels (in a Carrington longitude-latitude grid), 225  
 41 time frames, and output either 1 (emergence) or 0 (non-emergence).

42 We divide our data into a training+validation set (BMR observations within the first 10 months of every year between 1996  
 43 and 2011; Figure 2) and a test set (BMR observations on the last two months of every year between 1996 and 2011; Figure 2).  
 44 We exclusively use the training+validation set to train the CNN weights and optimize CNN hyper-parameters, and set aside the  
 45 test set to evaluate the performance of our algorithm under pseudo-operational conditions (i.e. with data that the CNN has  
 46 never seen). We make all our design decisions based on the training+validation set. The test set is only used to assess the final  
 47 performance shown this report.

48 For every model we augment each video with vertical and polarity flips resulting in an increase of the number of our training  
 49 and validation samples by a factor of 4. This is justified because the characteristics of magnetic flux emergence are not expected  
 50 to be different based on whether they occur in the northern or southern hemisphere of the Sun. Similarly, hemispheric polarity  
 51 orientation changes sign with every new solar cycle without this having any measurable impact on the properties of BMRs.

Months	Iterative Data Relabelling					Uncertainty Estimation			
	Performance Evaluation					Detection Threshold Definition			
	model1	model2	model3	model4	model5	model6	model7	model8	model9
Jan									
Feb					Validation				Validation
Mar									
Apr					Validation				
May									
Jun				Validation					
Jul			Validation						
Aug		Validation							
Sep									
Oct	Validation				Training set				Benchmarking set
Nov									
Dec					Test Set				

**Figure 2.** Data management used to build, train, validate and test different models for iterative relabelling, detection threshold estimation and performance evaluation. We split the data according to the months of the year to minimize the impact of the large temporal coherence present in solar data. The randomly-sampled splits typically used in other ML applications tend to overestimate performance in the solar case due to the fact that data close to each other in time can be generated by virtually the same structures.

We use a CNN<sup>15</sup> based on the Visual Geometry Group (VGG) architecture<sup>16</sup>. This architecture performs information abstraction with repeated convolution layers followed by a non-linear activation called ‘swish’ activation [ $\text{swish}(x) = \frac{x}{(1+e^{-x})}$ ]<sup>17</sup>. After each convolutional layer we perform max-pooling, gradually reducing the patch size and increasing the number of channels ending up in 1D vectors (Figure 1). Subsequently, the fully-connected layers at the end (Figure 1) generate the classification outcome with a sigmoid activation<sup>18</sup>. To regularize the model we add drop-outs<sup>19</sup> with a rate of 0.5 in the last two fully connected layers before the output layer. To optimize the network parameters during training we use a ‘stochastic gradient descent’<sup>20</sup> optimizer with a learning rate of  $10^{-4}$ , momentum of 0.9 and a batch size of 10.

For efficiently relabeling the data, we split the training+validation set (a total of 2032 videos) into 5 blocks of two months each and use the blocks to assemble five training+validation permutations to train 5 different models (models 1-5; Figure 2). This way, each one of these two-month blocks has an associated model where this block is not used to train CNN weights. This enables the unbiased evaluation of the classification of all videos and an assessment of the quality of the crude initial labeling.

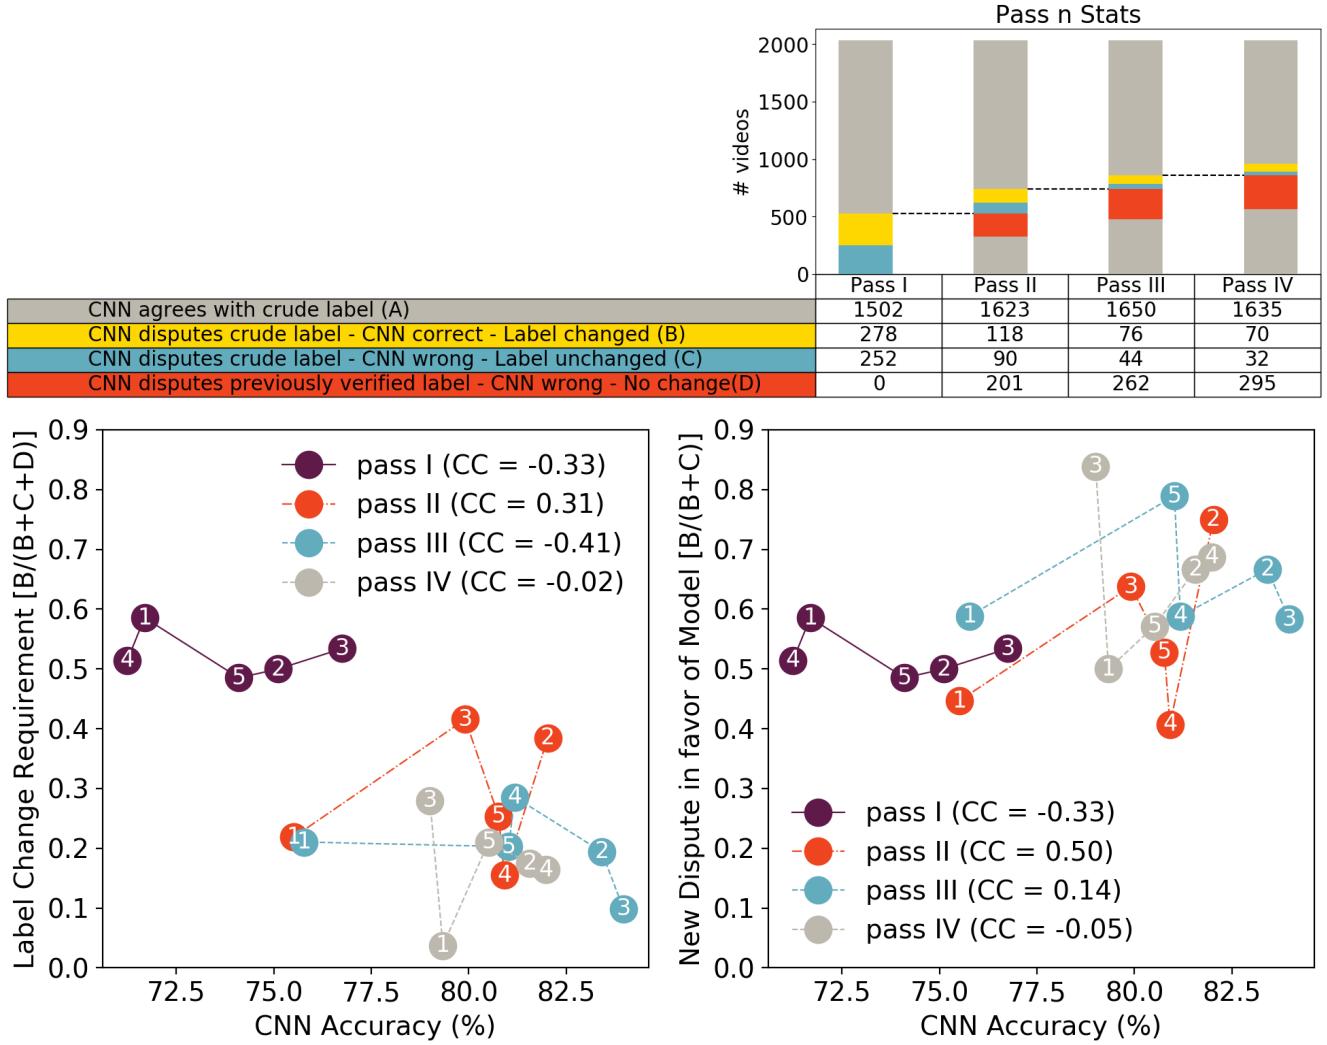
We train 5 models and evaluate them on their respective 5 non-overlapping validation blocks (Figure 2). Subsequently, we manually check and relabel (if necessary) the videos where CNN disagrees with the crude labeling (both false positives and false negatives). During each pass, we re-train all models with the verified data and repeat this verification and relabeling process until convergence (see detailed algorithm in ‘Methods’ section).

Figure 3 shows the relabelling process through a bar chart. We find that for higher passes an increasing fraction of the CNN labeling “mistakes” (false positives and negatives) are actually mistakes in the initial crude BARD labeling. After convergence, we manually check all the videos that the models classified successfully and find that ~90% of them had been correctly labelled through this process, even though we had only manually verified ~50% of them. This translates into a significant reduction of human effort when working with larger datasets.

We quantify the uncertainty in emergence/non-emergence classification using a different ensemble of the trained models (models 6-9; Figure 2). In this case we set aside months 9 and 10 as a common benchmark (‘benchmarking set’; Figure 2) and use the remaining four blocks to train and validate four permutation models like those described in the previous section. Additionally, we create two random realization per model by changing the random seed that determines the order in which the neural network sees the videos and the drop-out (random removal) of neuron connections.

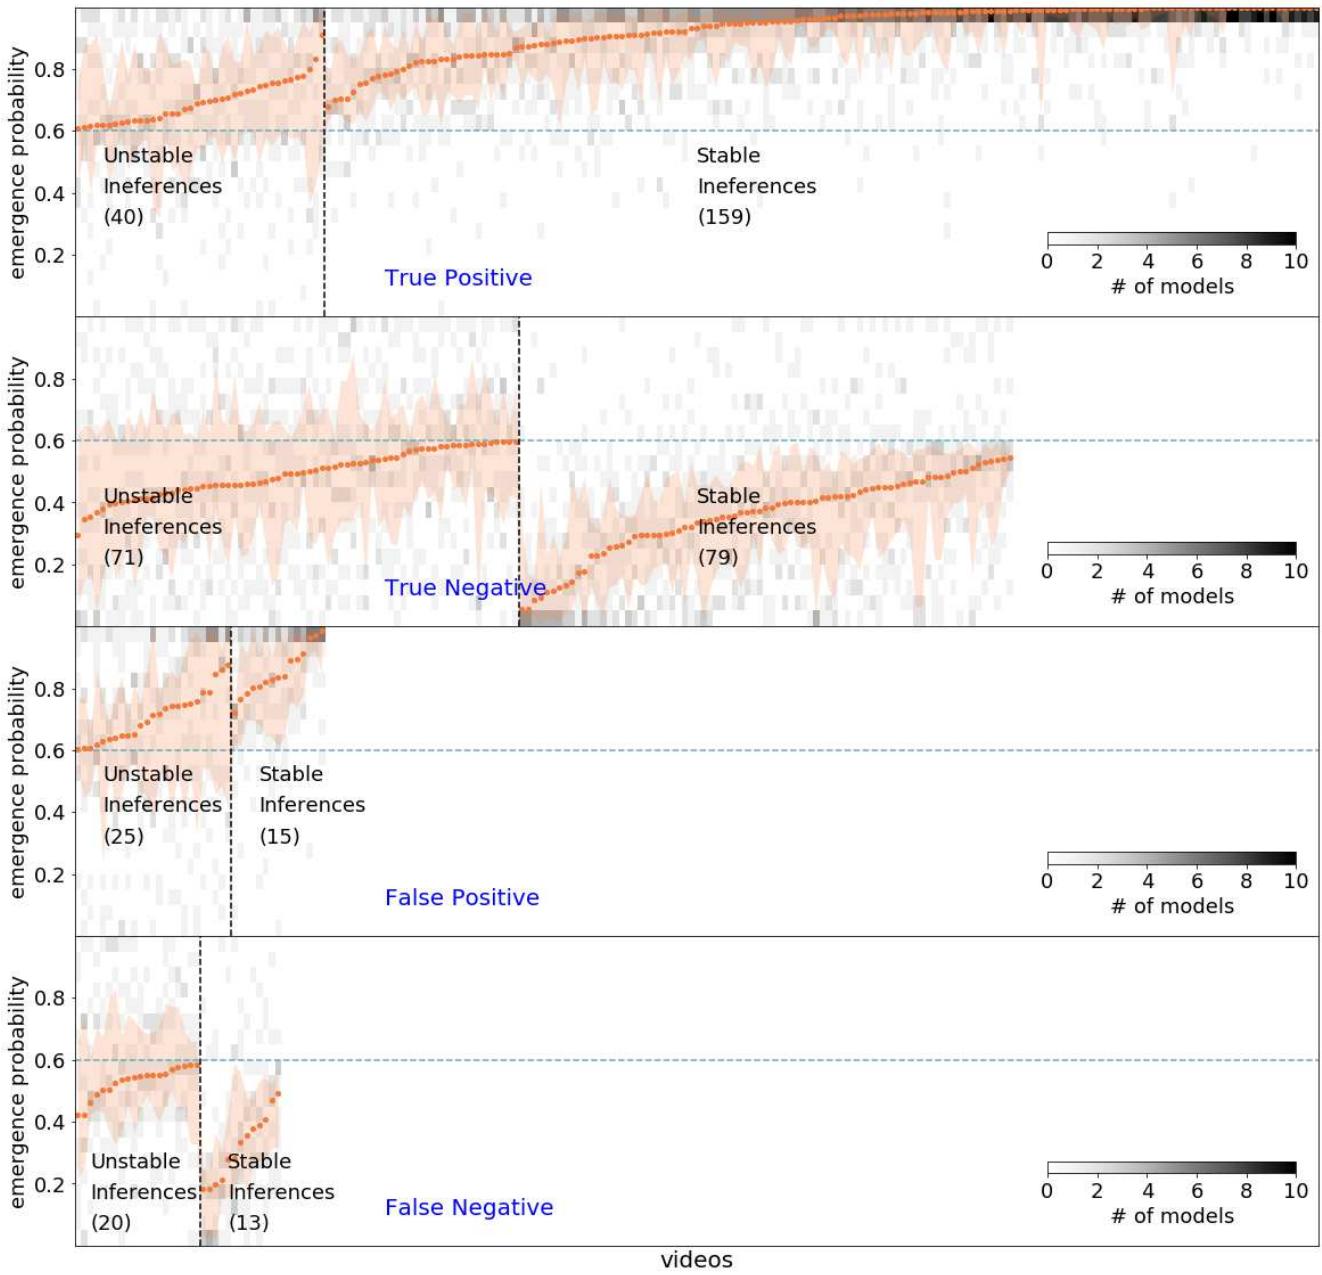
We use the median of the final sigmoid activation output of these 8 models (model 6-9  $\times$  2 random realizations) as the metric to classify each video. We find that an optimal threshold of 0.6 applied to the median of sigmoid model output yields the maximum accuracy. We derive this threshold from the intersection point of precision ( $\frac{\text{True Positive}}{\text{True Positive}+\text{False Positive}}$ ) and recall ( $\frac{\text{True Positive}}{\text{True Positive}+\text{False Negative}}$ ) curves.

Finally, we add one more random realization for each of models 1-5 and use the tuned 0.6 probability threshold (derived from models 6-9) to classify the ‘test set’ and assess performance under pseudo-operational conditions. Using model ensemble



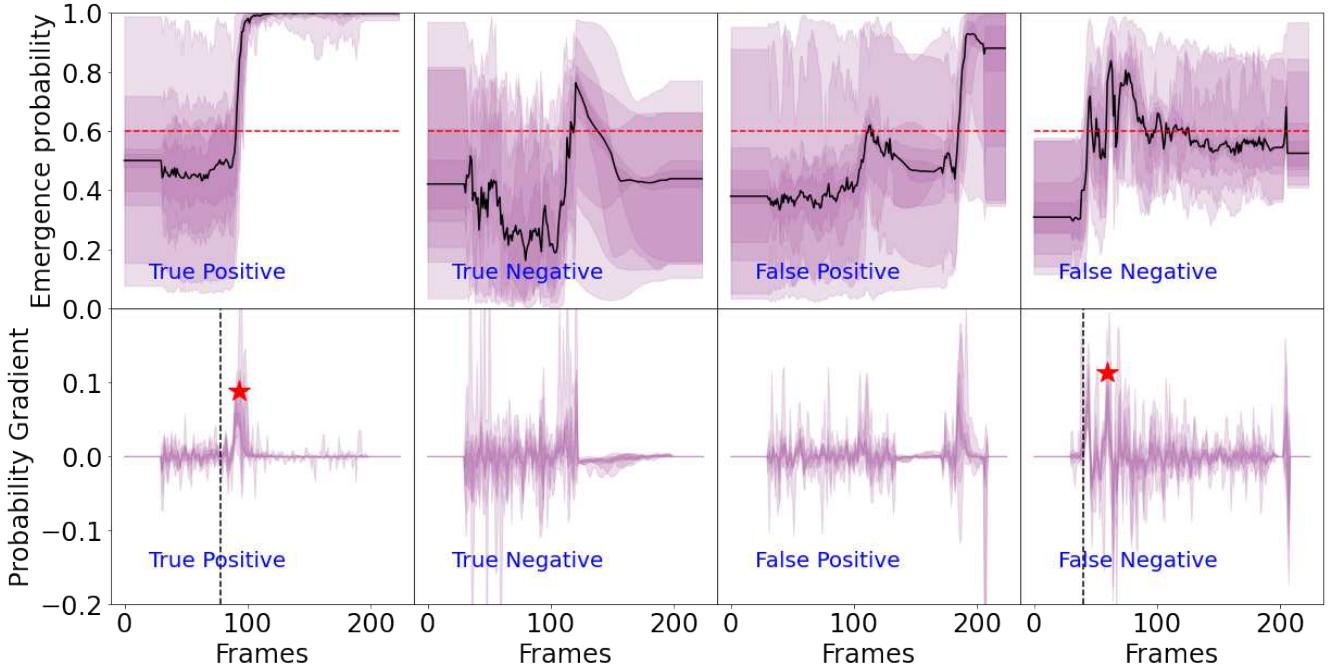
**Figure 3.** Sequence of iterative relabelling and convergence of performance. *Top panel* shows how the relabelling progresses through a total of 2032 videos (height of the bars) from Pass I to Pass IV. The videos are divided to 5 non-overlapping validation blocks (numbered circles in bottom panels) and are manually checked in a particular pass if the model (CNN) inference differs from the initial crude labelling and was not checked in prior passes. The red bars show the proportion of disagreements that have been manually verified in previous passes. The dashed black line shows the total proportion of videos that have been manually checked. The yellow (real mislabels) and blue (CNN mistakes) bars show newly discovered disagreements. The number of total disagreements reduces with each pass. The proportion of newly discovered disagreements that lead to a relabel increases with each pass. These trends result from labelled data makes more sense, allowing the CNN to better generalize. We quantify these trends as fractions  $\frac{B}{(B+C+D)}$ ,  $\frac{B}{(B+C)}$ , and plot them against CNN accuracy for different validation blocks in the bottom panels. *Bottom left panel* captures the trend in the bar chart and shows that the need of relabeling goes down from Pass I to Pass IV, because each pass is progressively downward in the plot. *Bottom right panel* shows that the models tend to capture the problems in the crude labeling better as we go higher in the passes because each pass is progressively upward in the plot. Both label change requirement and correctness of the CNN in disagreement become linearly independent of CNN accuracy for higher passes, as shown by the correlation coefficients.

83 median as the classification metric, we achieve 83% classification accuracy. Figure 4 shows model inferences on the ‘test  
 84 set’ as 2D histograms over videos and emergence probability, with the histogram frequency being the number of models that  
 85 output a certain probability value from sigmoid activation. We divide the 2D histograms into two regions, which we name  
 86 unstable and stable inferences for the purpose of uncertainty estimation in classification accuracy. Stable (unstable) inferences



**Figure 4.** Model ensemble used to estimate the classification accuracy on test set. Model inferences are depicted as 2D histograms. The x-axis indicates video number and the y-axis shows sigmoid output of the CNN-ensemble. The gray scale depicts the number of CNNs with a certain sigmoid probability output for a particular video. The median of all model outputs for a given video is marked with orange dots and the videos are arranged in ascending order of medians. Videos are classified as ‘emergence (positive)’ or ‘non-emergence (negative)’ depending on whether the median probability exceeds a threshold of 0.6 (dashed line) that is found as the intersection point of the precision and recall curves. The shaded regions enclose the 25<sup>th</sup> and 75<sup>th</sup> percentiles of the distribution of model outputs for each video. To estimate the uncertainty in classification accuracy, the ensemble inferences are divided into stable, unstable ones as shown in all the panels separated by a vertical dashed line. For stable (unstable) inferences, the probability threshold of 0.6 lies outside (inside) the ensemble 25<sup>th</sup>-75<sup>th</sup> percentile range.

<sup>87</sup> encompass the videos for which the probability threshold of 0.6 lies outside (inside) the model-ensemble’s 25<sup>th</sup>-75<sup>th</sup> percentile range. We define unstable classifications as those that will change if another percentile is used as reference. As observed in <sup>88</sup> Figure 4, the ensemble of CNNs is more confident in the classification of emergence than of non-emergence, i.e. there is a <sup>89</sup>



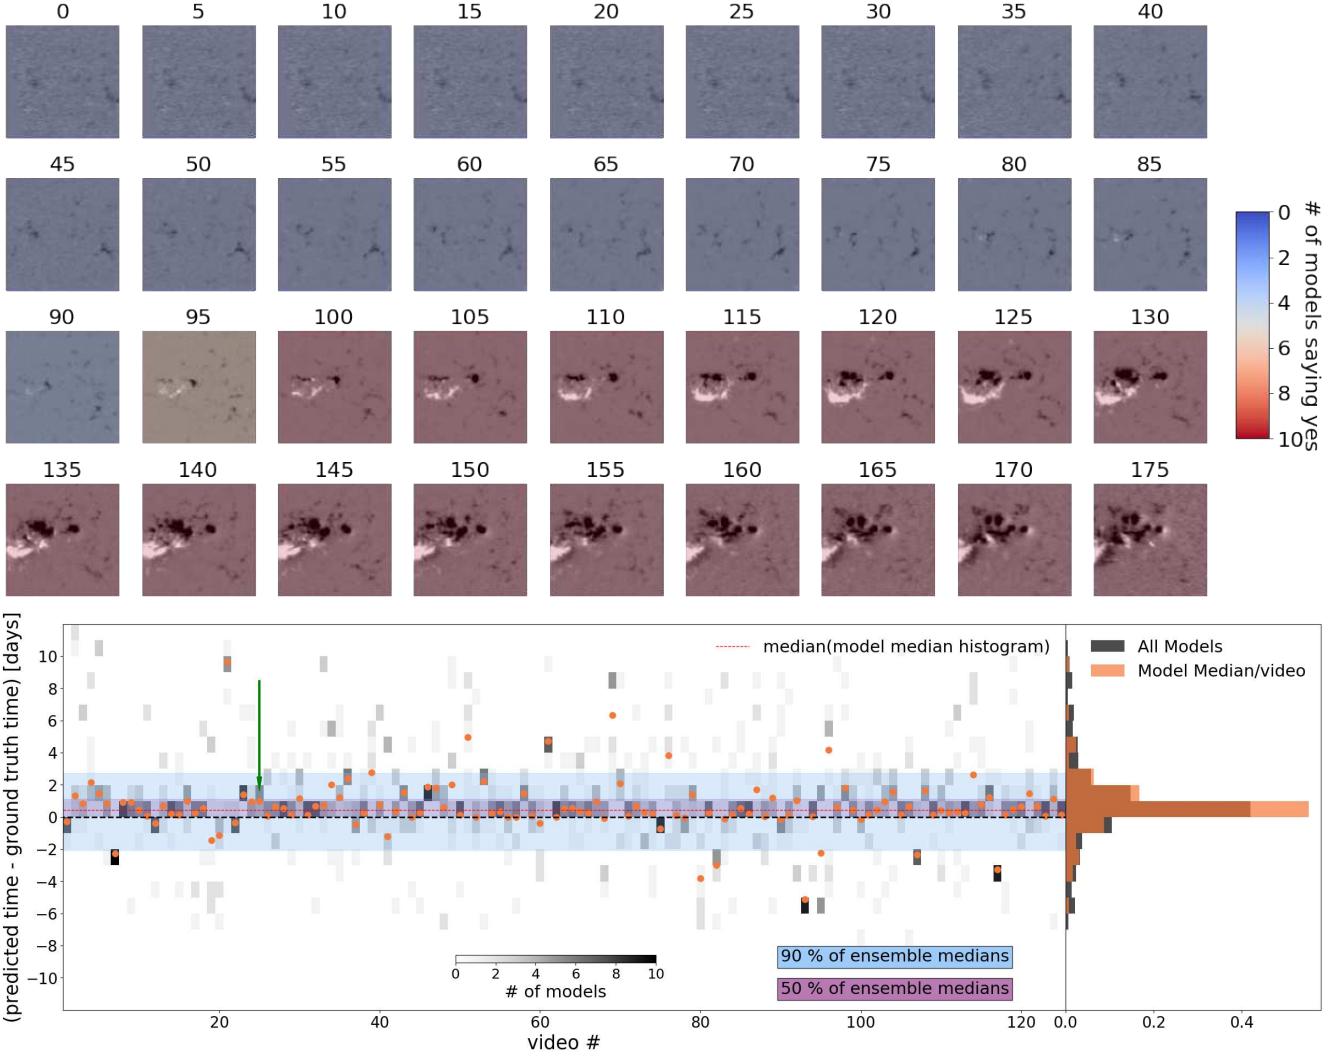
**Figure 5.** Identifying emergence epoch by frame stacking. *Top panels* show how the emergence probability changes for a video when frame truncation is gradually reduced to cover all the original frames and subsequently videos are classified as positive or negative. *Bottom panels* show the identification of emergence epoch (depicted by a red star) locating the frame of maximum emergence probability gradient. *Top left panel* depicts the same event as shown in Figure 1. The ground truth frames are marked by the vertical dashed lines. To show the central tendency of probability and probability gradient the region values for each frame number are sorted and regions between equidistant points from extrema are shaded.

90 larger (smaller) proportion of unstable true negatives (true positives). Because of this, performance is more sensitive to changes  
 91 in the classification of non-emergences than of emergences. We find that accuracy falls to 79% (71%) when the ensemble's 25<sup>th</sup>  
 92 (75<sup>th</sup>) percentile is used as the classification metric instead of the ensemble median. False positives and negatives have similar  
 93 ratios of stable vs. unstable videos. Thus, they do not seem to drive performance changes as much.

94 As a sanity check, we test the ability of the network to autonomously determine the time of each BMR emergence. For this,  
 95 we use all videos that the model ensemble classifies correctly. The idea is to identify what are the relevant frames that lead the  
 96 ensemble of models to classify a video as a BMR emergence. We mask-out portions of each video, after an arbitrary number of  
 97 frames, by repeating said frame until the end of the video. This is done iteratively for each video starting with the last frame,  
 98 then with the last two frames, then with the last three frames, then with the last  $n$  frames, and so on, until the entire video has  
 99 been masked out. We find there is a period in time where the majority of the models transition from classifying the video as an  
 100 emergence into a non-emergence. The frame when that happens coincides well with the moment where the BMR emerges on  
 101 the solar disk.

102 We quantify this by taking the time derivative of the sigmoid probability for each model in the ensemble and finding the  
 103 frame with maximum gradient. The median of the frame number for all models in the ensemble is what we determine to be  
 104 the frame of emergence. This process is shown in the top and bottom panels of Figure 5. The top-left and bottom-left panels  
 105 of Figure 5 depict the model-ensemble curves for one of the videos (frames shown in Figure 1). The ground-truth frame and  
 106 predicted frame locations are shown with vertical dashed line and red star respectively. We also show other sample curves  
 107 for false positives and true and false negatives. We note that the ensemble median at the end of the probability curve (i.e. no  
 108 masked frames) is what actually determines the classification (emergence/non-emergence), depending on whether the ensemble  
 109 median is greater or less than 0.6. Top panel of Figure 6 shows the time labeling of the BMR emergence, where frame color  
 110 indicates the number of models classifying the video as an emergence given that all subsequent frames have been masked-out.

111 The bottom panel of Figure 6, shows the accuracy of our emergence frame detection. On the left we show a 2D histogram  
 112 with the difference between the estimated and real time of emergence for all test set videos that the model-ensemble correctly  
 113 classifies as BMR emergence. The median of all model ensembles is +0.4 days. This means that using the median of the  
 114 location of maximum probability gradient tends to result in an estimation that is slightly behind the observed emergence frame.



**Figure 6.** Emergence epoch identification using an ensemble of models. *Top panel* shows every 5th frame of the ‘bipolar emergence’ video also shown in Figures 1 & 5. Each frame is color-coded according to the number of models that classify the video as ‘emergence’ until that frame. *Bottom panel* shows a 2D histogram representing the joint distribution of the ‘bipolar emergence’ videos and predicted emergence time w.r.t. ground truth using model ensemble. The gray level represents the number of models predicting within a time bin (1 day) for a particular video. The green arrow points to the location of the video in top panel in the 2D histogram. Median time for over the models for each video is depicted by orange dot. The ground truth reference is depicted by the black dashed horizontal line. The histograms on the right show occurrence of prediction over all models (black) and model ensemble medians per video (orange). The median of the orange histogram is depicted by the dashed orange horizontal line. It clearly shows that the models have a tendency to detect the emergence little later ( $\approx 0.4$  day) than the actual initiation (also seen in Figure 5). The purple and light blue shaded region show the time range within which 25<sup>th</sup>-75<sup>th</sup> percentile and 5<sup>th</sup>-95<sup>th</sup> percentile of predictions lie respectively.

115 The 5<sup>th</sup> percentile is -2.1 days, the 25<sup>th</sup> percentile +0.1 days, the 75<sup>th</sup> percentile +1.1 days, and the 95<sup>th</sup> percentile +2.7 days.  
 116 This means that 50% (90%) of emergence detections are within [0.1,1.1]([-2.1, 2.7]) days of the observed emergence. For  
 117 reference, a  $15^\circ \times 15^\circ$  patch is visible on disk for  $\sim 14$  days.

118 We note that we have not fine-tuned or optimized in any way what are the optimal detection thresholds that would  
 119 to maximize accuracy of frame detection. This is beyond the scope of this paper and needs to be done carefully on the  
 120 training+validation set before it can be tested on a pseudo-operational setting with the test set. Instead, our goal is to showcase  
 121 the versatility and potential of CNNs. We want to highlight how training them for something as simple as a ‘yes/no’ question

122 can be easily re-purposed for a much more sophisticated question like a 'yes/no and, if yes, when?'.

## 123 Discussion

124 We show that a deep learning model can be harnessed to refine the initially crude labelling of the dataset used for training  
125 the model. We achieve this by partitioning the training dataset into blocks and training several models with complementary  
126 validation sets. We find that, as we progress through this iterative process, any false positive and false negatives reported by  
127 the CNN are increasingly likely to be mislabels. This iterative process lessens the manual effort needed for painstakingly  
128 labeling data by 50%, which is one of the main obstacles of deep learning applications to classification on large astronomical  
129 databases. We factored in the associate uncertainty of using a relatively small dataset by training an ensemble of models and  
130 calculating the median and range of possible classification outcomes. Using the ensemble median yields a classification  
131 accuracy of 83% of BMR flux emergences vs. non-emergences. Choosing the 25<sup>th</sup> (75<sup>th</sup>) percentile, instead of the median,  
132 translates into a performance reduction to 79% (71%). This asymmetry is a consequence of the CNN ensemble being more  
133 confident about classifying emergences vs. non-emergences.

134 Even though our model was solely trained as an emergence classifier, we show that it can be re-purposed to also detect  
135 the time of emergence by progressively freezing video frames until the emergence event is masked. We see this as evidence  
136 that our model is being able to learn, abstract, and generalize the characteristics that make a BMR flux emergence. This has  
137 interesting implications for the labeling, and classification of dynamical astrophysical events in which the exact time of the  
138 event is unknown. This implies that it may be possible to use the deep learning model for the prediction of emergence by only  
139 looking at the early evolution prior to flux emergence. Early detection of magnetic flux emergence, if possible before there are  
140 signatures visible to a human observer, is a holy grail in space weather forecasting. Coupled with a early-warning observatory  
141 placed at the L1 Lagrangian point, it would significantly increase our readiness and ability to mitigate its impact.

## 142 Methods

143 **BARD.** The Bipolar Active Region Detection (BARD) catalog<sup>14</sup> uses a semi-automatic segmentation algorithm, coupled with  
144 human supervision to detect and track BMRs as they emerge and/or rotate in and out of view at a cadence of one observation  
145 per day.

146 **VGG.** This is a CNN architecture available in two varieties- VGG16 and VGG19 named after the depth of layers<sup>16</sup>. The  
147 convolutional layers uses  $3 \times 3$  kernels. The volume is controlled by maxpooling as the depth increases. The network ends with  
148 two fully-connected layers. It is widely used for natural image classification. Our CNN architecture is based on VGG with  
149 input being videos instead of images.

150 **Sigmoid Activation.** The sigmoid function<sup>18</sup> ( $\text{sigmoid}(x) = \frac{1}{1+e^{-x}}$ ) outputs a number within 0-1 and is typically used to  
151 represent classification probability. We use sigmoid activation at the end to evaluate the probability of a video to be an  
152 'emergence'.

153 **Swish Activation.** Swish<sup>17</sup> a smooth activation function developed by Google to replace the widely used activation function  
154 ReLU ( $\max(0, x)$ ). Swish activation is represented as  $\text{swish}(x) = x \cdot \text{sigmoid}(x)$ . We use 'swish' activation as the non-linearity  
155 after every convolution layer.

156 **Relabelling Algorithm.** The detailed steps of our iterative relabelling algorithm are shown below-

- 161 1. Divide crudely labelled data into 5 train-validation combination where  $\text{Months}_{\text{val}} = \{2m - 1, 2m\}$  and  $\text{Months}_{\text{train}} =$   
162  $\{1, 2, \dots, 10\} \setminus \text{Months}_{\text{val}}$   $\forall m \in \{1, 2, 3, 4, 5\}$
- 163 2. Train the model (CNN) for each combination until validation accuracy reaches maximum.
- 164 3. Using the trained model inference identify false-positives ( $fp$ ) and false-negatives ( $fn$ ).
- 165 4. Look at each  $fp$ ,  $fn$  and identify model mistakes ( $E$ )
- 166 5. Calculate the mistakes of crude label  $|fp| + |fn| - |E|$  and change those labels
- 167 6. Calculate the ratio  $R = \frac{|E|}{|fp| + |fn| - |E|}$
- 168 7.  $n = 1$

169        8. Data<sub>Pass[n]</sub> = relabelled data  
 170        9. **While** R> 0.5 **do**  
 171            (a) Divide Data<sub>Pass[n]</sub> into 5 train-validation combination where Months<sub>val</sub> = {2m – 1, 2m} and Months<sub>train</sub> = {1, 2, .., 10}  
 172            \Months<sub>val</sub>  $\forall m \in \{1, 2, 3, 4, 5\}$   
 173            (b) Train the model for each combination until validation accuracy reaches maximum.  
 174            (c) Using the trained model inference identify new (not seen in prior passes) false-positives (*fp*) and false-negatives  
 175            (*fn*).  
 176            (d) Look at each *fp*, *fn* and identify model mistakes (*E*)  
 177            (e) Calculate the mistakes of crude label  $|fp| + |fn| - |E|$  and change those labels  
 178            (f) Calculate the ratio  $R = \frac{|E|}{|fp| + |fn| - |E|}$   
 179            (g) n = n + 1  
 180            (h) Data<sub>Pass[n]</sub> = relabelled data

181        10. **return** Data<sub>Pass[n]</sub>, model

## 182 Acknowledgements

183 This research was funded by NASA grants 80NSSC19M0165 and 80NSSC18K0671.

## 184 Author contributions statement

185 S. Chatterjee and A. Muñoz-Jaramillo planned the experiments and wrote the paper.  
 186 S. Chatterjee set up and ran the experiments.  
 187 A. Muñoz-Jaramillo provided the list of events that was analyzed.  
 188 D. A. Lamb assembled the video sequences used in this work and helped edit the manuscript.

## 189 Competing interests

190 The authors declare no competing interests.

## 191 References

- Zhang, Y. & Zhao, Y. Astronomy in the Big Data Era. *Data Sci. J.* **14**, 11 (2015). DOI 10.5334/dsj-2015-011.
- Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In Pereira, F., Burges, C. J. C., Bottou, L. & Weinberger, K. Q. (eds.) *Advances in Neural Information Processing Systems 25*, 1097–1105 (Curran Associates, Inc., 2012).
- Settles, B. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison (2009).
- Dubey, G., van der Holst, B. & Poedts, S. The initiation of coronal mass ejections by magnetic flux emergence. *A & A* **459**, 927–934 (2006). DOI 10.1051/0004-6361:20054719.
- Zhang, Y., Zhang, M. & Zhang, H. On the Relationship between Flux Emergence and CME Initiation. *Sol. Phys.* **250**, 75–88 (2008). DOI 10.1007/s11207-008-9150-7.
- Rycroft, M. J. *Space Weather and Hazards to Application Satellites*, 1175–1193 (Springer New York, New York, NY, 2013). DOI 10.1007/978-1-4419-7671-0-78.
- DeForest, C. E., Hagenaar, H. J., Lamb, D. A., Parnell, C. E. & Welsch, B. T. Solar magnetic tracking. i. software comparison and recommended practices. *The Astrophys. J.* **666**, 576–587 (2007). DOI 10.1086/518994.
- Lamb, D. A., DeForest, C. E., Hagenaar, H. J., Parnell, C. E. & Welsch, B. T. Solar magnetic tracking. II. the apparent unipolar origin of quiet-sun flux. *The Astrophys. J.* **674**, 520–529 (2008). DOI 10.1086/524372.

- 208 9. Iida, Y., Hagenaar, H. J. & Yokoyama, T. Detection of Flux Emergence, Splitting, Merging, and Cancellation of Network  
209 Field. I. Splitting and Merging. **752**, 149 (2012). DOI 10.1088/0004-637X/752/2/149. [1204.5261](#).
- 210 10. Iida, Y., Hagenaar, H. J. & Yokoyama, T. Detection of Flux Emergence, Splitting, Merging, and Cancellation of Network  
211 Fields. II. Apparent Unipolar Flux Change and Cancellation. *The Astrophys. J.* **814**, 134 (2015). DOI 10.1088/0004-  
212 637x/814/2/134.
- 213 11. Iida, Yusuke. Tracking of magnetic flux concentrations over a five-day observation, and an insight into surface magnetic  
214 flux transport. *J. Space Weather. Space Clim.* **6**, A27 (2016). DOI 10.1051/swsc/2016022.
- 215 12. Jiang, H. *et al.* Identifying and tracking solar magnetic flux elements with deep learning. *The Astrophys. J. Suppl. Ser.* **250**,  
216 5 (2020). DOI 10.3847/1538-4365/aba4aa.
- 217 13. Scherrer, P. H. *et al.* The Solar Oscillations Investigation - Michelson Doppler Imager. *Sol. Phys.* **162**, 129–188 (1995).  
218 DOI 10.1007/BF00733429.
- 219 14. Muñoz-Jaramillo, A. *et al.* The best of both worlds: Using automatic detection and limited human supervision to create a  
220 homogenous magnetic catalog spanning four solar cycles. In *2016 IEEE International Conference on Big Data (Big Data)*,  
221 3194–3203 (2016).
- 222 15. LeCun, Y., Bottou, L. & Bengio, Y. *Object Recognition with Gradient-Based Learning*, 319–345 (Springer  
223 Berlin Heidelberg, Berlin, Heidelberg, 1999). DOI 10.1007/3-540-46805-6-19.
- 224 16. Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv e-prints*  
225 arXiv:1409.1556 (2014). [1409.1556](#).
- 226 17. Ramachandran, P., Zoph, B. & Le, Q. V. Searching for Activation Functions. *arXiv e-prints* arXiv:1710.05941 (2017).  
227 [1710.05941](#).
- 228 18. Han, J. & Moraga, C. The influence of the sigmoid function parameters on the speed of backpropagation learning. In Mira,  
229 J. & Sandoval, F. (eds.) *From Natural to Artificial Neural Computation*, 195–201 (Springer Berlin Heidelberg, Berlin,  
230 Heidelberg, 1995).
- 231 19. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: A simple way to prevent neural  
232 networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
- 233 20. Ruder, S. An overview of gradient descent optimization algorithms. *CoRR* **abs/1609.04747** (2016). [1609.04747](#).

# Figures

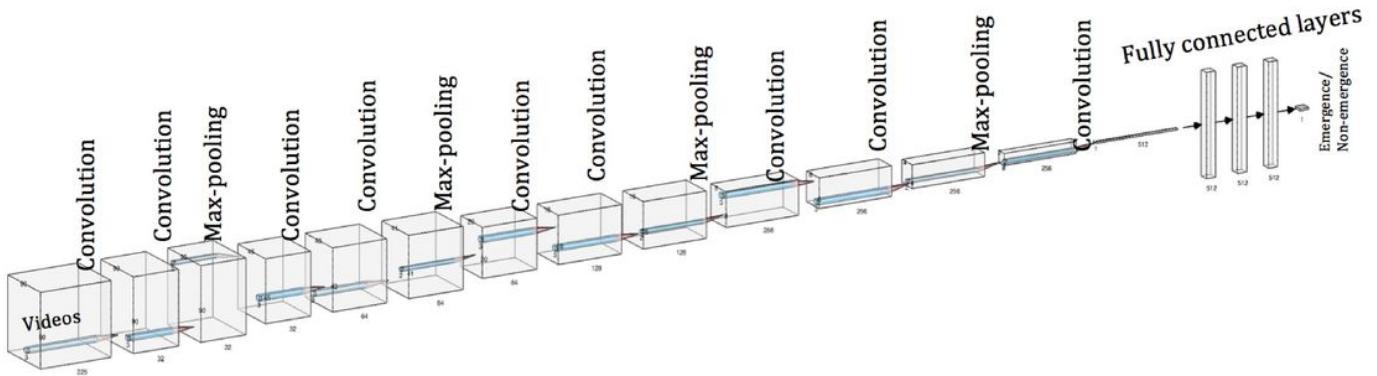
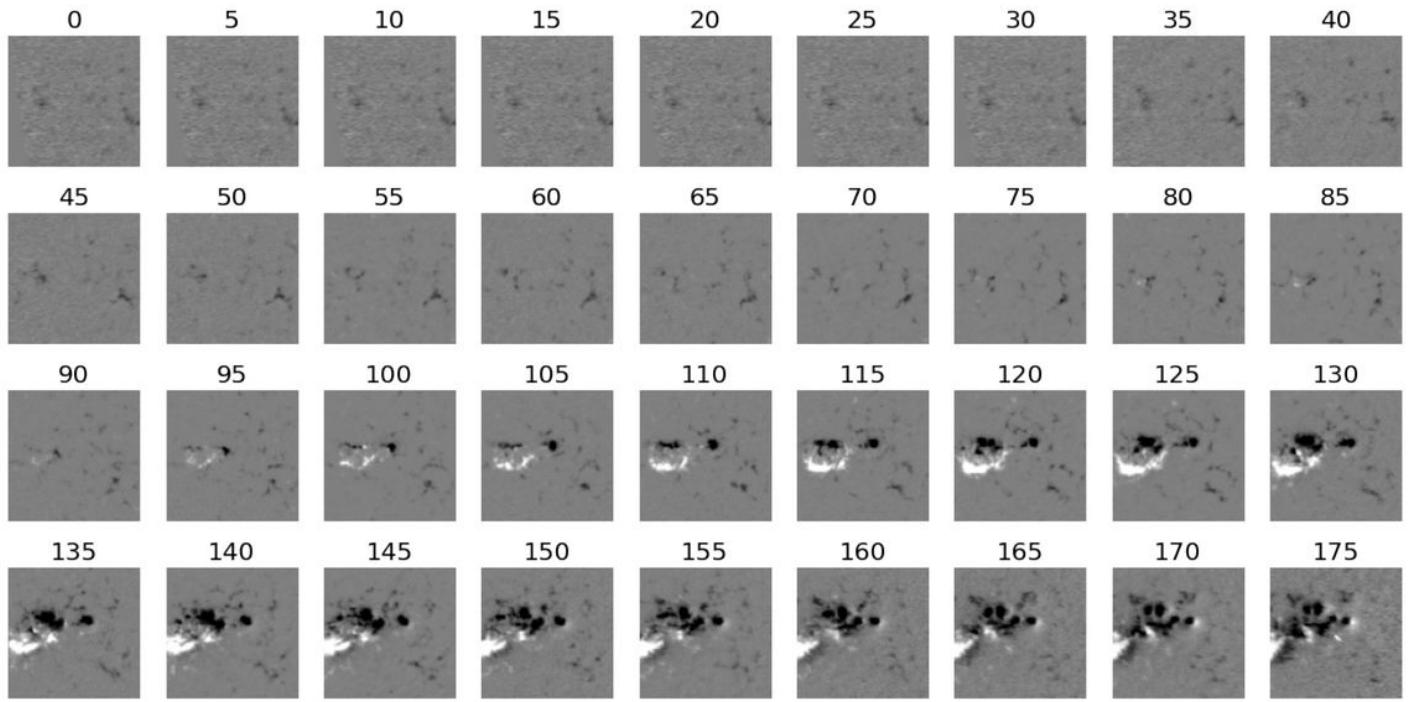
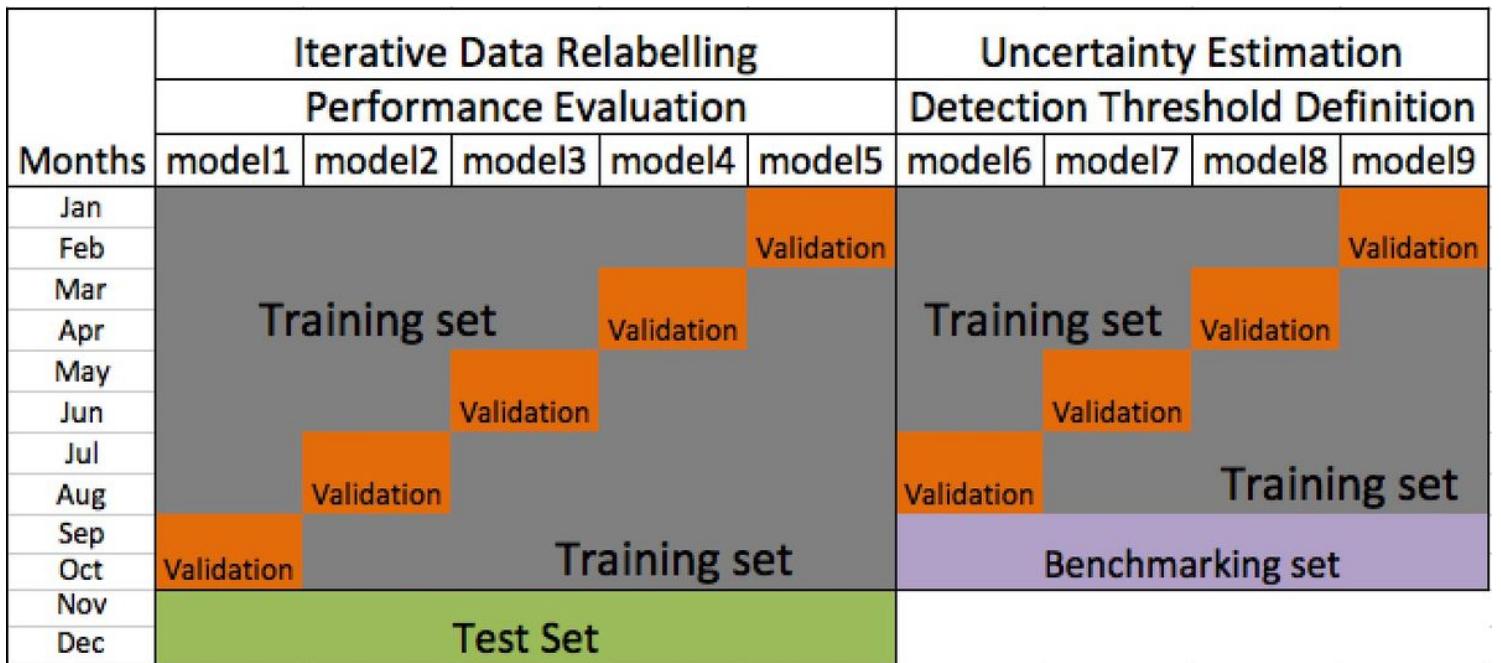


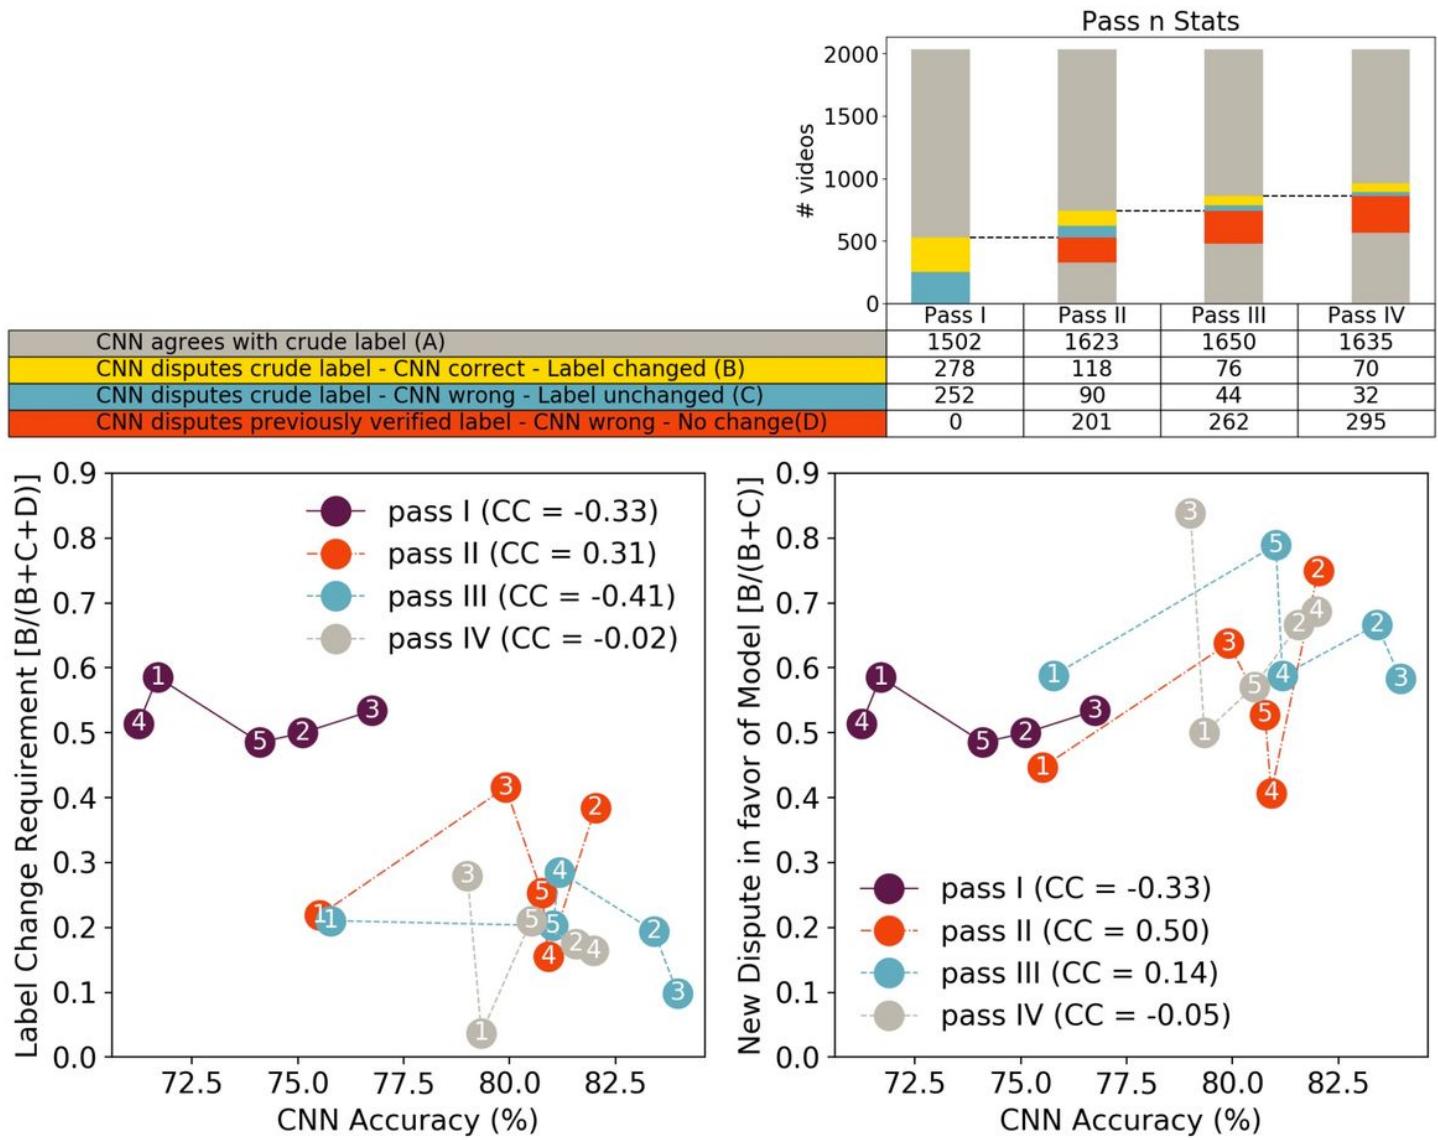
Figure 1

Dataset and model architecture. Top panel shows every 5th frame of an emergence video starting on November 29, 1998 17:36 hours (UT). Bottom panel shows adapted deep learning model for classifying videos into 'emergence' or 'non-emergence' classes. We use repeated convolution, max-pooling and fully connected layers.



**Figure 2**

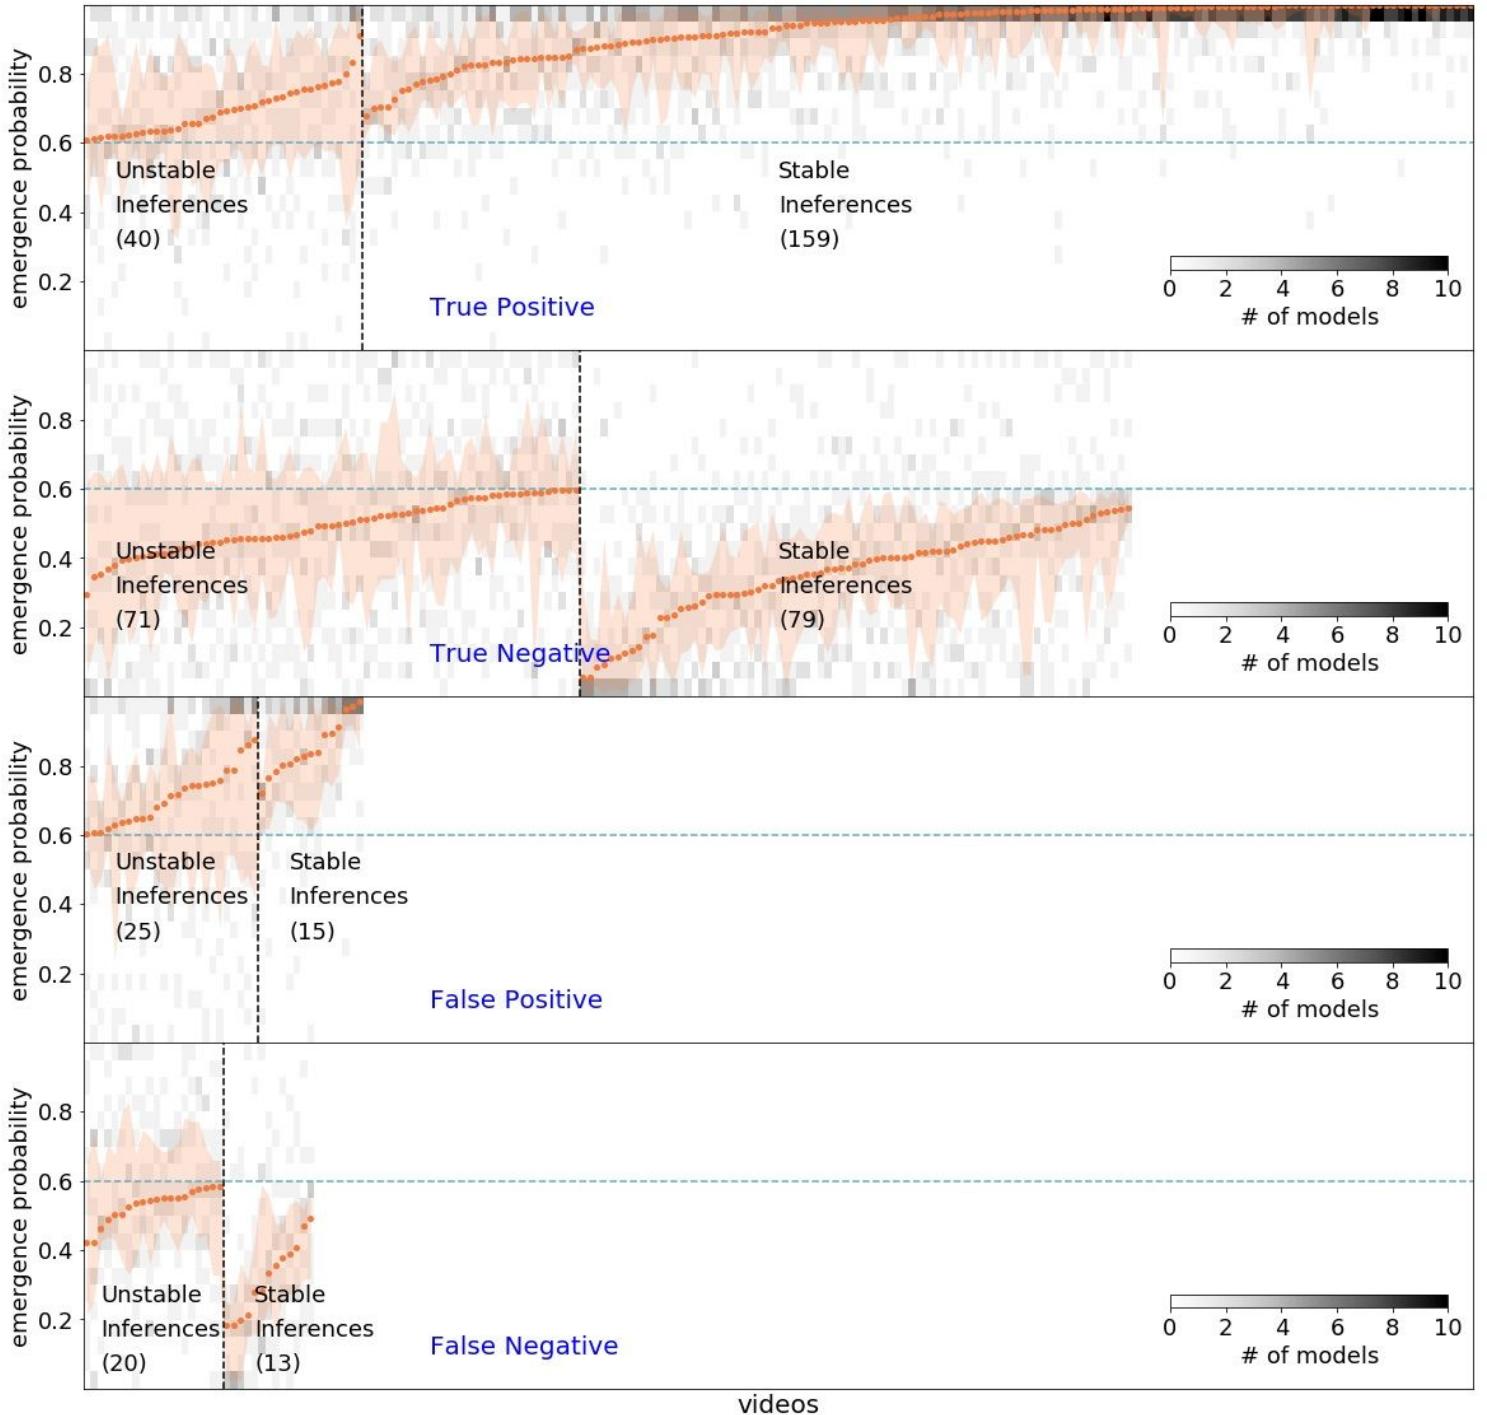
Data management used to build, train, validate and test different models for iterative relabelling, detection threshold estimation and performance evaluation. We split the data according to the months of the year to minimize the impact of the large temporal coherence present in solar data. The randomly-sampled splits typically used in other ML applications tend to overestimate performance in the solar case due to the fact that data close to each other in time can be generated by virtually the same structures.



**Figure 3**

Sequence of iterative relabelling and convergence of performance. Top panel shows how the relabelling progresses through a total of 2032 videos (height of the bars) from Pass I to Pass IV. The videos are divided to 5 non-overlapping validation blocks (numbered circles in bottom panels) and are manually checked in a particular pass if the model (CNN) inference differs from the initial crude labelling and was not checked in prior passes. The red bars show the proportion of disagreements that have been manually verified in previous passes. The dashed black line shows the total proportion of videos that have been manually checked. The yellow (real mislabels) and blue (CNN mistakes) bars show newly discovered disagreements. The number of total disagreements reduces with each pass. The proportion of newly discovered disagreements that lead to a relabel increases with each pass. These trends result from labelled data makes more sense, allowing the CNN to better generalize. We quantify these trends as fractions  $B/(B+C+D)$ ;  $B/(B+C)$ , and plot them against CNN accuracy for different validation blocks in the bottom panels. Bottom left panel captures the trend in the bar chart and shows that the need of relabeling goes down from Pass I to Pass IV, because each pass is progressively downward in the plot. Bottom right

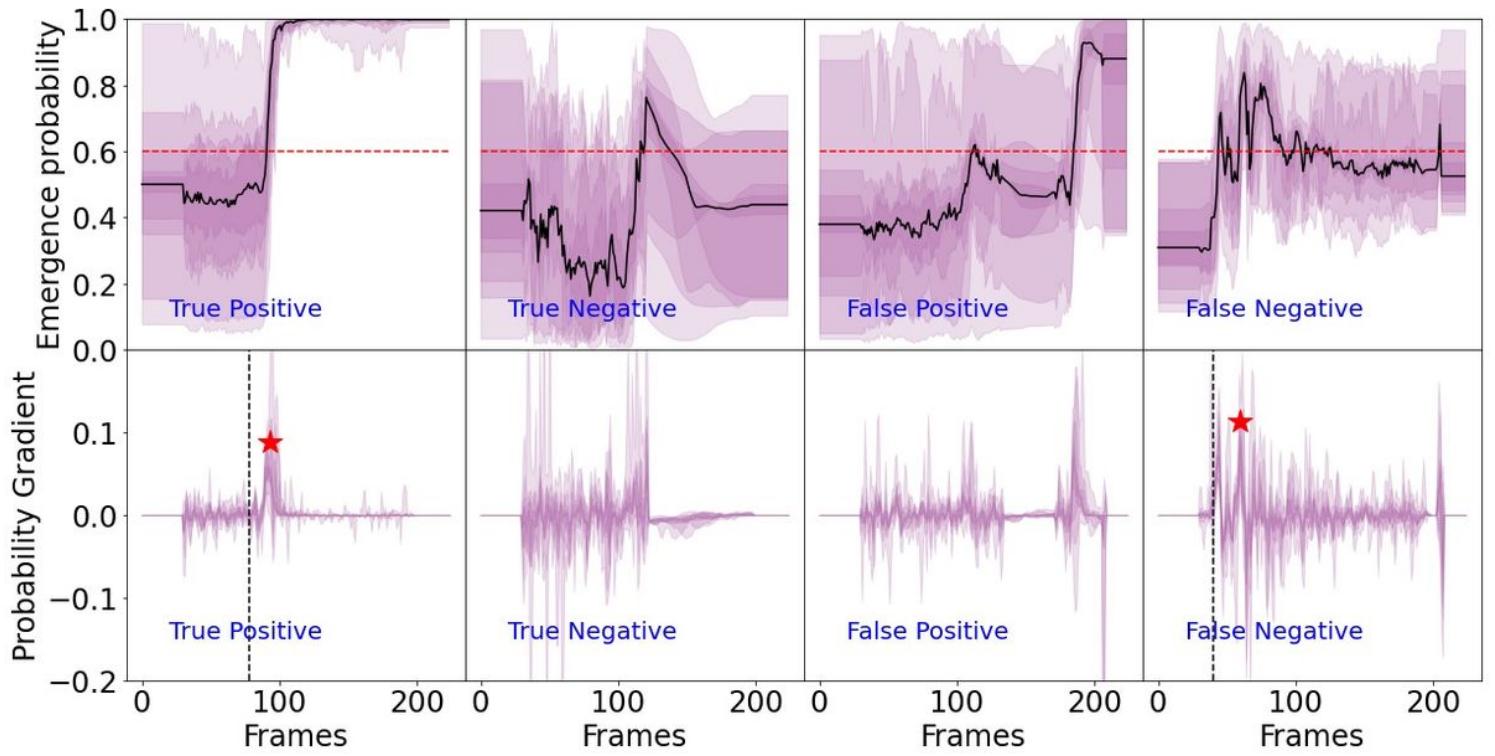
panel shows that the models tend to capture the problems in the crude labeling better as we go higher in the passes because each pass is progressively upward in the plot. Both label change requirement and correctness of the CNN in disagreement become linearly independent of CNN accuracy for higher passes, as shown by the correlation coefficients.



**Figure 4**

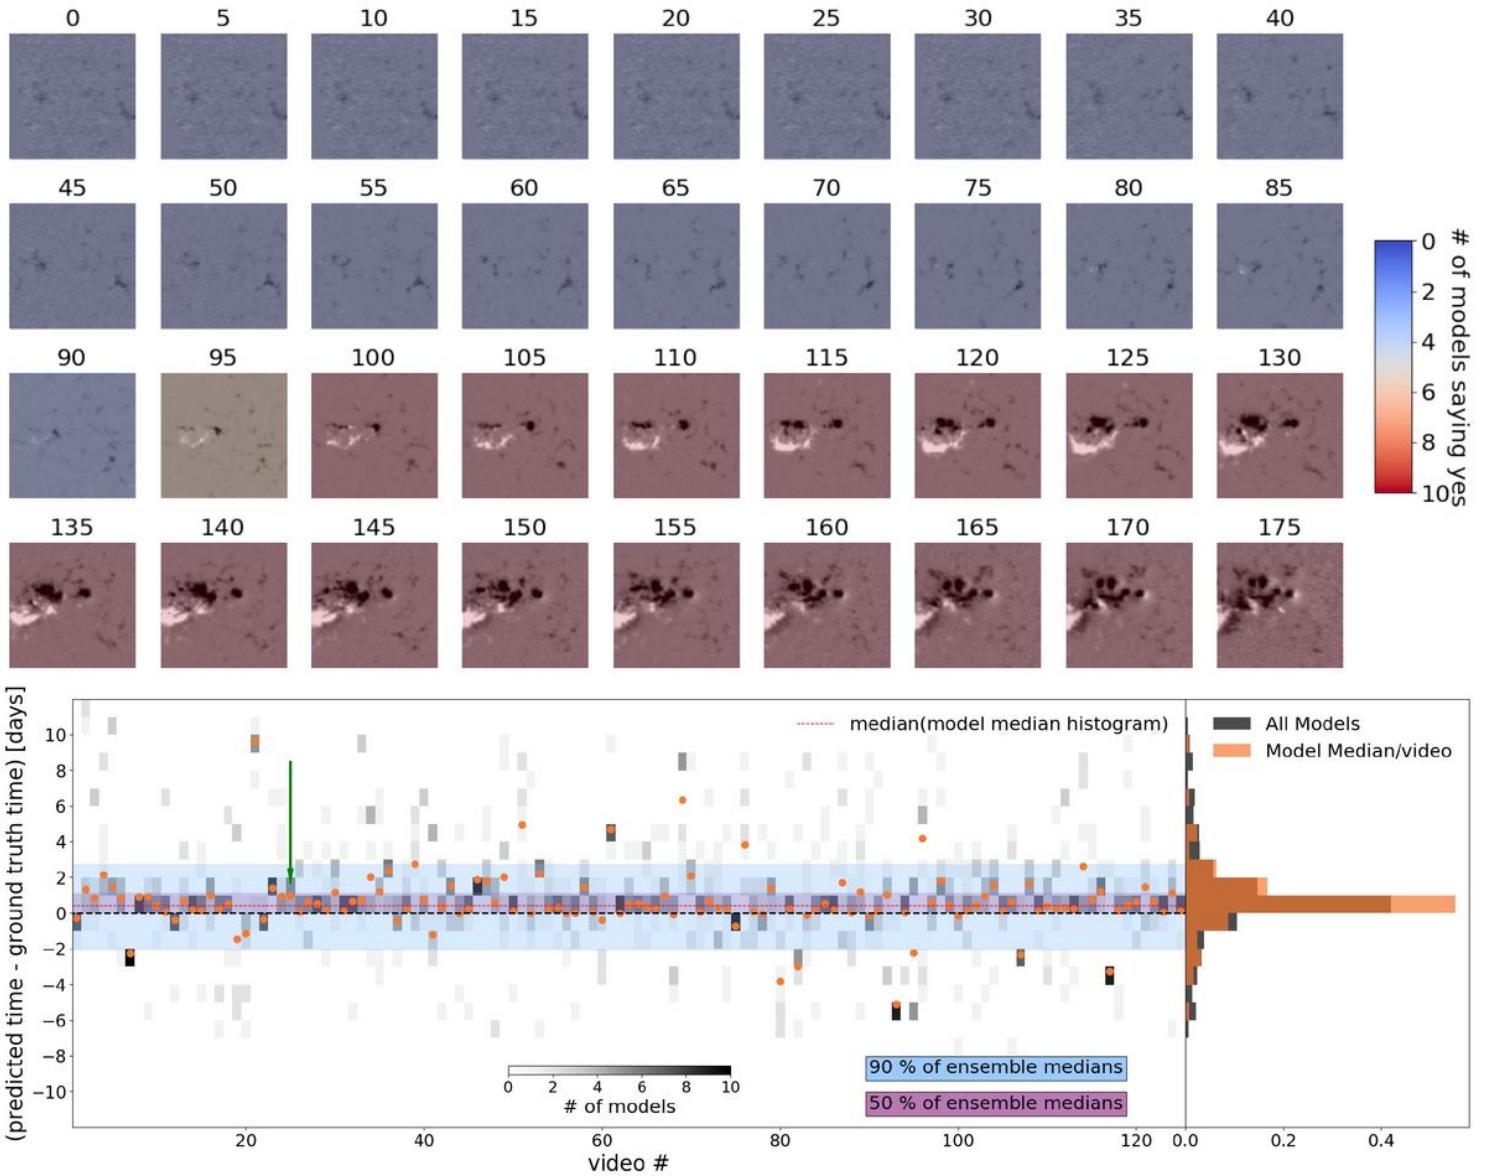
Model ensemble used to estimate the classification accuracy on test set. Model inferences are depicted as 2D histograms. The x-axis indicates video number and the y-axis shows sigmoid output of the CNN-ensemble. The gray scale depicts the number of CNNs with a certain sigmoid probability output for a

particular video. The median of all model outputs for a given video is marked with orange dots and the videos are arranged in ascending order of medians. Videos are classified as ‘emergence (positive)’ or ‘non-emergence (negative)’ depending on whether the median probability exceeds a threshold of 0.6 (dashed line) that is found as the intersection point of the precision and recall curves. The shaded regions enclose the 25th and 75th percentiles of the distribution of model outputs for each video. To estimate the uncertainty in classification accuracy, the ensemble inferences are divided into stable, unstable ones as shown in all the panels separated by a vertical dashed line. For stable (unstable) inferences, the probability threshold of 0.6 lies outside (inside) the ensemble 25th-75th percentile range.



**Figure 5**

Identifying emergence epoch by frame stacking. Top panels show how the emergence probability changes for a video when frame truncation is gradually reduced to cover all the original frames and subsequently videos are classified as positive or negative. Bottom panels show the identification of emergence epoch (depicted by a red star) locating the frame of maximum emergence probability gradient. Top left panel depicts the same event as shown in Figure 1. The ground truth frames are marked by the vertical dashed lines. To show the central tendency of probability and probability gradient the region values for each frame number are sorted and regions between equidistant points from extrema are shaded.



**Figure 6**

Emergence epoch identification using an ensemble of models. Top panel shows every 5th frame of the 'bipolar emergence' video also shown in Figures 1 & 5. Each frame is color-coded according to the number of models that classify the video as 'emergence' until that frame. Bottom panel shows a 2D histogram representing the joint distribution of the 'bipolar emergence' videos and predicted emergence time w.r.t. ground truth using model ensemble. The gray level represents the number of models predicting within a time bin (1 day) for a particular video. The green arrow points to the location of the video in top panel in the 2D histogram. Median time for over the models for each video is depicted by orange dot. The ground truth reference is depicted by the black dashed horizontal line. The histograms on the right show occurrence of prediction over all models (black) and model ensemble medians per video (orange). The median of the orange histogram is depicted by the dashed orange horizontal line. It clearly shows that the models have a tendency to detect the emergence little later (0.4 day) than the actual initiation (also seen in Figure 5). The purple and light blue shaded region show the time range within which 25th-75th percentile and 5th-95th percentile of predictions lie respectively.