

1 **Attention-based machine vision models and techniques**
2 **for solar wind speed forecasting using solar EUV**
3 **images**

4 **Edward J. E. Brown**^{1,2,*}, **Filip Svoboda**^{1,*}, **Nigel P. Meredith**², **Nicholas Lane**
5 ^{1,3}, **Richard B. Horne**²

6 ¹Department of Computer Science and Technology, University of Cambridge, Cambridge, UK

7 ²British Antarctic Survey, Natural Environment Research Council, Cambridge, UK

8 ³Samsung AI Center, Cambridge, UK

9 *equal contribution

10 **Key Points:**

- 11 • Attention-based machine vision models and methodological enhancements are de-
12 veloped to improve solar wind speed forecasts from solar images
- 13 • Attention-based architectures outperform convolutional models, motivating their
14 use in future studies and production systems
- 15 • The models perform best in the declining phase of the solar cycle when activity
16 is driven by coronal holes

17 **Abstract**

18 Extreme ultraviolet images taken by the Atmospheric Imaging Assembly on board the
19 Solar Dynamics Observatory make it possible to use deep vision techniques to forecast
20 solar wind speed - a difficult, high-impact, and unsolved problem. At a four day time
21 horizon, this study uses attention-based models and a set of methodological improvements
22 to deliver an 11.1% lower RMSE error and a 17.4% higher prediction correlation com-
23 pared to the previous work testing on the period from 2010 to 2018. Our analysis shows
24 that attention-based models combined with our pipeline consistently outperform con-
25 volutional alternatives. Our model has learned relationships between coronal holes' char-
26 acteristics and the speed of their associated high speed streams, agreeing with empir-
27 ical results. Our study finds a strong dependence of our best model on the position in
28 the solar cycle, with the best performance occurring in the declining phase.

Plain language summary

Solar images contain rich information that can be used to forecast conditions at Earth. This study develops a robust methodology for processing solar images and trains machine learning models that can use them to predict the solar wind speed. Combined, these deliver a very significant 17.4% improvement in the correlation between the prediction and the ground truth over previous works. The models perform better during the quieter, declining phase of the solar cycle when the solar activity is driven by coronal holes. Finally, the trained models learn properties of coronal holes that agree with prior empirical studies.

1 Introduction

The solar wind is a stream of charged particles that is emitted from the upper atmosphere of the Sun. The speed, density, temperature and the magnitude and direction of the associated magnetic field of the solar wind are constantly varying affecting the way in which it ultimately interacts with the Earth's magnetosphere. High speed solar wind streams (HSS) emanating from coronal holes are particularly effective at coupling with the Earth's magnetosphere. The weak storms they produce tend to have long-lasting recovery phases which often result in prolonged and enhanced substorm activity (Tsurutani et al., 1995; Meredith et al., 2011). This results in repeated injections of suprathermal electrons into the inner magnetosphere and significant increases in the fluxes of relativistic electrons in the outer radiation belt, increasing the risk to satellites via surface charging and internal charging respectively (e.g., Borovsky and Denton (2006)). Indeed, it has been suggested that satellites at geostationary orbit are more likely to be at risk from an extreme HSS-driven storm than a Carrington type event (Horne et al., 2018). Furthermore, prolonged and enhanced substorm activity associated with HSS-driven storms results in increased thermospheric densities and satellite drag (Chen et al., 2012). Consequently, accurately forecasting the solar wind speed associated with coronal holes is very important for our modern society.

Coronal holes are large dark areas on the Sun as seen in extreme ultraviolet (EUV) and soft X-ray images (Cranmer, 2009). They are regions of open magnetic field and cooler plasma, leading to the production of high speed solar wind streams. Coronal holes are long-lasting features that can persist from one solar rotation to the next, giving rise to a 27 day periodicity in the arrival of HSS at Earth. The occurrence rate of coronal holes peaks during the declining phase of the solar cycle (Burlaga & Lepping, 1977) and high speed streams observed at Earth during these intervals tend to be coronal-hole driven. The distribution of speeds in high speed streams associated with coronal holes ranges from 400 to 800 kms^{-1} (Kilpua et al., 2017). While these streams do not result in ma-

65 jor geomagnetic storms (Richardson et al., 2006), they have extensive recovery phases,
66 typically lasting from 5–10 days, and, as a result, may deposit more energy in the mag-
67 netosphere than larger storms (Kozyra et al., 2006; Turner et al., 2006).

68 Coronal holes are not the only source of high speed solar wind at Earth. Coronal mass
69 ejections (CMEs) also cause high speed solar wind, although not all CMEs are associ-
70 ated with high solar wind speeds (Kilpua et al., 2017). CME’s are large explosions on
71 the Sun that hurl vast amounts of plasma into space. The occurrence rate of CMEs peaks
72 at solar maximum (St. Cyr et al., 2000) so that most periods of high solar wind speed
73 observed during these periods tend to be CME-driven. The distribution of speeds in in-
74 terplanetary coronal mass ejections (ICMEs) and sheath regions associated with CMEs
75 on the Sun ranges from 250 to 950 kms^{-1} (Kilpua et al., 2017). Unlike coronal holes,
76 CMEs are not associated with long lasting features on the Sun. In contrast they are best
77 observed in coronagraph images where they appear as expanding shells of material.

78 In this study we build a machine learning model to use solar images to forecast the so-
79 lar wind speed at Earth. This technique is expected to perform best when there are as-
80 sociated visible features on the Sun. The method is thus expected to work well for coro-
81 nal holes, which are large features on the solar disk. In contrast, coronal mass ejections
82 are barely noticeable within EUV images and so the ML model would not be expected
83 to work well for these events.

84 The field of machine learning has built a lot of momentum over the last 10 years. This
85 has largely been the result of improvements in algorithmic capability, availability of data,
86 funding and hardware. Not to be overlooked though is the creation of field benchmarks
87 like ImageNet (Deng et al., 2009) and open-source software such as PyTorch (Paszke et
88 al., 2019) which dramatically shortened the development cycle in the field and greatly
89 increased its standardization.

90 Deep (Machine) Learning excels where rich data exists in large quantities, because mod-
91 els with deep structures and therefore many parameters need to consume richly varied
92 data sources to build complex internal representations of the data generating system.
93 This is the essence of deep learning. Recently, curated solar image datasets have been
94 created such as the SDO ML dataset (Galvez et al., 2019) which contains images of the
95 Sun taken at various EUV wavelengths. These data allow the rapid application of ma-
96 chine learning algorithms to consume solar images.

97 In this paper we use the EUV images taken by the Solar Dynamics Observatory (SDO)
98 using the Atmospheric Image Assembly (AIA) (Lemen et al., 2011) to forecast the so-

lar wind speed at the Lagrangian L1 point. We present results for forecasting at a four day lag from a single 211 Å image - but this forecast could be used for any lag up to four days. We also explore the model’s learned behaviour by examining relationships between the peak solar wind speed and the coronal hole area and intensity. Previous works and the datasets are presented in Sections 2 and 3 respectively. In Section 4, we discuss our general methodology and model architectures. Our results are presented and discussed in 5. Finally, our conclusions are summarised in Section 6.

2 Previous Works

The works of Wintoft and Lundstedt (1997) and Wintoft and Lundstedt (1999) were the first to use neural networks to forecast the solar wind speed. These are small, so-called fully connected, models that could learn non-linear relationships between a limited set of pre-computed feature inputs, such as the flux tube expansion factor, and the solar wind speed. More recently, similar studies were performed by D. D. Liu et al. (2011), Yang et al. (2018), Chandorkar et al. (2019), and Bailey et al. (2021) using similar non-image-based inputs to the models, albeit with more advanced models than the earlier works.

Upendran et al. (2020) was the first study aiming to forecast solar wind speed from solar EUV images using deep learning techniques. The work uses images from both 193 and 211 Angstrom wavelengths to forecast the solar wind speed at a one day resolution. Upendran uses GoogleNet (Szegedy et al., 2014), trained on the ImageNet dataset (Deng et al., 2009), as a feature extractor for each image. The extracted per-image features are then passed into an LSTM Recurrent Neural Network (Hochreiter & Schmidhuber, 1997) to produce the predicted solar wind speed. The study achieves a best performing model at a lag of 3 days and a history of 4 days, with a correlation of 0.55 and an RMSE 80.28 km/s. This study will build on this insightful initial work.

Next, Raju and Das (2021) proposed a smaller three-layer convolutional feature extractor, which they train on the 193 Angstrom wavelength solar EUV images. Their method targets a subtly different task than that of Upendran et al. (2020). While Upendran et al. (2020) forecast future solar wind speeds based on images at a fixed distance in the present, Raju and Das (2021) backcast current solar wind speed based on flexible-lag past images. Specifically, Raju and Das (2021) use the current solar wind speed to infer which past image was likely to have caused the recorded solar wind speed, and then pass this image into their model with the expectation that the model will be able to correctly reconstruct the observed solar wind speed. The key difference between the two approaches is that in the forecasting setup the model needs to use the image information to both infer the speed of the caused solar wind and make a judgement on whether the particle stream at the given speed will be geo-effective. In the Raju and Das (2021) backcast-

135 ing setup, the observed features are guaranteed to have been geo-effective, as the image
 136 was chosen based on this criterion, and thus the model needs to infer only the speed, not
 137 the geo-effectiveness too. Its task is thus made easier. When used in prediction it is as-
 138 sumed that the predicted solar wind speed will be used to infer the time when it will ar-
 139 rive at Earth. The difference becomes clearer when the models are to be deployed as so-
 140 lar wind speed predictors. Under the forecasting setup, today’s images can be used to
 141 produce the predicted solar wind speed 4 days from now. In contrast, under the back-
 142 casting setup not all time stamps would receive a prediction. Indeed, the inference pro-
 143 cess by which images are paired with time stamps does not guarantee a unique predic-
 144 tion for each time stamp, and so some time stamps can be expected to receive multiple
 145 solar wind speed predictions, while others would get none. Furthermore, any error in the
 146 speed prediction will be significantly magnified when the time offset is made dependent
 147 on this prediction. Thus this model is not comparable to Upendran et al. (2020). Nev-
 148 ertheless, they provide results for a model specially trained at a fixed 4 day forecast hori-
 149 zon (their Table 4), with the year 2018 held out as a test set. They report 78.3 km/s RMSE
 150 and a prediction correlation of 0.55. This would be comparable to Upendran et al. (2020),
 151 except they provide no results for 2018 alone. Their test results are from across multi-
 152 ple years. Therefore, our study will compare to Upendran et al. (2020) for dates across
 153 an 8.5 year range and then run a separate training run to compare to Raju and Das (2021)’s
 154 fixed 4-day model, just evaluating on the year 2018.

155 **3 Data**

156 **3.1 Solar Images**

157 The image dataset consists of EUV images from NASA’s Solar Dynamics Observatory
 158 (SDO) taken by the Atmospheric Imaging Assembly (AIA) (Lemen et al., 2011) that have
 159 been processed by performing various instrumental corrections, downsampled to usable
 160 spatial and temporal resolutions and synchronised both spatially and temporally to form
 161 the SDOML dataset (Galvez et al., 2019). The resulting dataset contains 8 and a half
 162 years of images every 6 minutes from June 2010 to December 2018. These images are
 163 monochromatic and the pixel values represent the intensity of light. This study uses the
 164 EUV images at 211 Angstroms.

165 **3.2 Solar Wind Speed**

166 The solar wind speed data are taken from the OMNIWeb service. Specifically, we use
 167 the solar wind speed, measured in km/s, at a 1 minute time resolution for the period of
 168 the SDOML dataset. The data come from WIND and the Advanced Composition Ex-
 169 plorer (ACE) spacecraft, both positioned at the L1 point, about 1.5 million km from Earth.

170 The solar wind speed is highly auto-correlated with itself over hourly time periods and
 171 is still at 0.7 after 1 day. By four days, the correlation has dropped to negligible amounts.
 172 Notably, at 27 days, there is a spike in the auto-correlation. This is because the Sun has
 173 a synodic rotation period of approximately 27 days and some longer lasting features, such
 174 as coronal holes, come around again causing similar solar wind speed conditions at L1.
 175 This auto-correlation is important since it has implications for which images are included
 176 in training and test sets due to their dependence on each-other. This is further discussed
 177 in Section 4.1.7.

178 **4 Methodology**

179 **4.1 Methodological Improvements**

180 Here we discuss changes in our methodology to the only previous work, (Upendran et
 181 al., 2020), covering all the date ranges available from the SDO ML dataset.

182 ***4.1.1 Image pre-processing***

183 The EUV images at their provided resolution are too large to practically process on stan-
 184 dard computing hardware. Previous works elected to down-sample the full 512 by 512
 185 pixel image to 224 by 224 by max pooling. Instead, we take a 300 by 300 pixel square
 186 who’s corners are approximately at the edges of the solar disk, and then down sample
 187 this cropped image to the desired 224 by 224 image size. This results in lower loss of in-
 188 formation content in the relevant section of the Sun because 1) the cropped solar poles
 189 are unlikely to be geo-effective, 2) the cropped features at the western limb take about
 190 7 days to be geo-effective and so are outside of the max 4 day forecasting horizon, 3) this
 191 allowed us to down-sample the central, relevant, portion of the image less aggressively.
 192 Figure 1 shows an example of our cropping technique.

193 Regarding scaling the cropped image images, the same method as used in Upendran et
 194 al. (2020) is employed by clipping the pixels to have values between a minimum of 25
 195 and a maximum of 2500 and taking the natural logarithm. However, after this we rely
 196 on a batchnorm layer to learn an optimal scaling, as opposed to fixing it (further detailed
 197 in Section 4.2).

198 ***4.1.2 Sampling frequency***

199 We replace the previously used daily sampling resolution with a 30 minute schedule, be-
 200 cause solar wind speeds can change significantly even on a 30 minute time scale.

201 ***4.1.3 Carrington rotation***

202 The Sun rotates on average every 27.28 days as viewed from Earth, this is one Carring-
 203 ton rotation(Ridpath, 2012). As such, the solar features that affected the solar wind speed

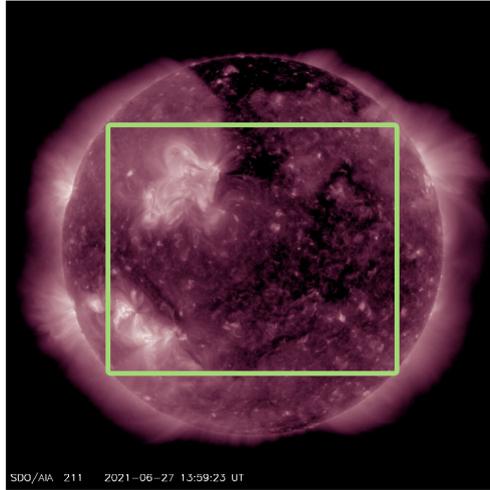


Figure 1: SDO AIA 211Å image taken on 2021-06-27 (Lemen et al., 2011)

204 at a given point come back approximately 27 days later and produce similar effects. Thus,
 205 the solar wind speed is also auto-correlated at the Carrington rotation periodicity with
 206 a value of 0.42 at 27 days. As this value is available to all forecasters operating at lower
 207 than 27 days forecast horizon, it should be used as an input to our models.

208 ***4.1.4 North-south augmentation***

209 We augment the dataset by randomly flipping the training images north to south, as fea-
 210 tures, such as coronal holes, produce a similar increase in solar wind speed regardless of
 211 which side of the solar equator they are on. Although it is not claimed these are valid
 212 physical suns.

213 ***4.1.5 Single image versus sequence***

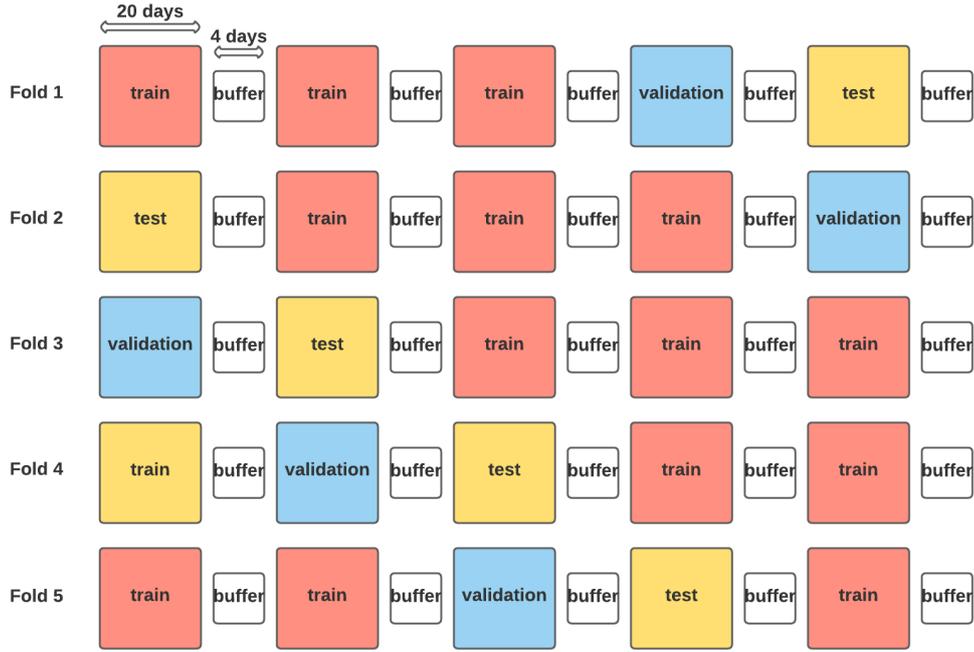
214 The previous work relies on a convolutional feature extractor pre-trained on ImageNet
 215 in combination with an LSTM cell and a fully connected layer (Upendran et al., 2020).
 216 Up to 4 images were sequentially passed through the convolutions. Separate for each im-
 217 age, the model’s activations at multiple layers were extracted, concatenated, and passed
 218 into the LSTM as individual time steps. The convolutions remained parametrized by the
 219 weights obtained on ImageNet and only the other layers’ parameters were trained. The
 220 high auto-correlation of solar images is likely to, again, exaggerate the model’s multi-
 221 collinearity in hidden features while providing little additional context. Thus we replaced
 222 the LSTM feeding into a fully connected output layer with two consecutive fully con-
 223 nected layers.

224 **4.1.6 Feature extractor re-training**

225 This study will use pre-trained vision models at the core of the model architecture (see
226 Section 4.2 for more details). Rather than to use the fixed pre-trained ImageNet weights,
227 the model will be initialised with these weights but they will not be fixed. This we be-
228 lieve to be strictly necessary due to the wide gap between the EUV and the ImageNet
229 datasets.

230 **4.1.7 Training, validation and test sets**

231 For this study, 5-fold cross-validation is employed to evaluate the models. Solar wind speed
232 is auto-correlated up to a period of about 4 days. For the period of June 2010 to Decem-
233 ber 2018, the auto-correlation is as high as 0.70 at one day. This means that if times-
234 tamps are too close to each-other between training, validation and test sets, it is not a
235 fair reflection of the performance of a model, since the Sun has not changed much in for
236 example 30 minutes. Furthermore, this will mean that the model overfits on the valida-
237 tion sets, meaning they will not generalise as well. In order to create more independent
238 training and test sets, a method similar to that used in Upendran et al. (2020) is em-
239 ployed whereby the timestamps from 2010-2018 are split into chunks of 20 days. How-
240 ever, a buffer period of 4 days between each chunk is thrown out to ensure the indepen-
241 dence of the training, validation and test sets. It is noted that this throws out approx-
242 imately one fifth of all the data. However, this is justified to ensure the independence
243 of datasets while also covering as many parts of the solar cycle as possible. These chunks
244 are then put into training, validation and test buckets. This process is repeated 5 times
245 to ensure that each 20 day chunk serves a turn in the training set 3 times, the valida-
246 tion set once and the test set once. This creates 5 folds of training, validation and test.
247 For each fold, a model is trained on the training set and evaluated on the validation set
248 for 100 epochs (1 epoch is a full pass over the data). The model is saved every epoch.
249 The version of the model that performs best on the validation set is the final model. This
250 final model is then applied and evaluated on the unseen test set. Figure 2a shows the
251 training sets in orange, the validation sets in blue and the test sets in yellow. White buffer
252 sets of 4 days are included between the 20 day chunks.



(a) 5-fold cross validation with buffer data thrown out. Pattern is repeated across the May 2010 to December 2018 range.



(b) Dataset split with 2018 as hold-out test set for comparison with Raju and Das (2021)

Figure 2: Training, Validation and Test sets

253 **4.2 Model Architectures**

254 For this study, the architectures for the different models will follow the format in Fig-
 255 ure 3. The image will pass through a batch norm layer that will rescale it. Then it is passed
 256 into the candidate architecture, be it a CNN or a vision transformer. The outputs from
 257 this model as well as the solar wind speed from one Carrington rotation ago are then passed
 258 into two final consecutive non-linear projections that produce the model’s solar wind speed
 259 prediction.

260 In all cases the models are trained in their entirety on the EUV data. That is, after their
 261 parameters are initialized using either random, or when available, pre-set weights the al-
 262 gorithm iteratively updates them with the goal of incrementally decreasing the mean squared
 263 error of its prediction.

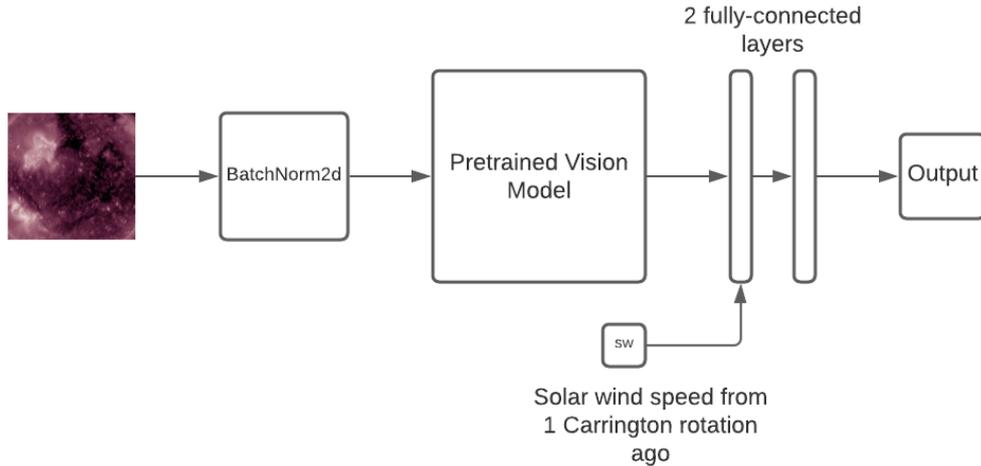


Figure 3: The Solar model architecture

264 **4.2.1 Benchmark CNN-based models**

265 In general every deep model can be seen as a layered composition of non-linear projec-
 266 tions, each forming a separate layer. Model inputs, solar images in our case, can be seen
 267 as the zero-th layer, while, model outputs, the predicted solar wind speed, can be treated
 268 as the last layer. Each layer in between is a non-linear projection that receives inputs
 269 from the preceding layer, and that outputs its value to the next layer. Commonly, sev-
 270 eral layers are grouped into modules and used as a type of meta-layer. Modern archi-
 271 tectures are defined by the features that build on and expand this basic structure.

272 Previous work used convolutional models in the forecasting of solar wind, (Upendran et
 273 al. (2020); Raju and Das (2021)). These models are designed to process images, each of
 274 which has three dimensions - the height, the width, and the number of channels. A stan-
 275 dard colour image has 3 channels: red, green, and blue. Convolutions are operations that
 276 split the image into a grid of patches and then use a three dimensional kernel to com-
 277 pute weighted averages per each patch. The same kernel is used on each patch and the
 278 averages it produces become the pixel values the layer outputs. Multiple kernels may be

279 employed, in which case their outputs are treated as separate channels of the outputted
280 image.

281 **GoogleNet**, also known as InceptionNet v1, is the convolutional architecture at the heart
282 of Upendran et al. (2020)’s work. It is a convolutional architecture that replaces layers
283 with modules. Each module computes several, rather than just one convolution. These
284 are computed in parallel, and are meant to complement each other. The desired effect
285 is to make the model’s computation more parallelizable, thus faster, while improving the
286 model’s ability to fit complex patters in the data (Szegedy et al., 2014).

287 **InceptionNet v2** is a second generation and a refinement of the GoogleNet. The ar-
288 chitecture builds on GoogleNet’s inception modules by decomposing their convolutions
289 serially. Specifically, more computationally expensive, that is larger-kernel convolutions,
290 are replaced by a series of much cheaper smaller-kernel convolutions carried out one af-
291 ter the other. The desired effect is to make the working set of this algorithm smaller, while
292 further improving the model’s capacity, i.e. its ability to fit complex data patterns (Szegedy
293 et al., 2016).

294 **ResNet** is a predecessor of GoogleNet. ResNet’s modules consist of two consecutive con-
295 volutions, and a so called residual connection. The residual connection is a bypass that
296 circumvents the two convolutions. In effect this results in a block that outputs both its
297 convolution’s output as well as the original inputs to the block. This trick helps to prop-
298 agate the training gradients through the network, mitigating the vanishing gradient prob-
299 lem. The architecture was the first one to breach the 20 layer depth ceiling (He et al.,
300 2016).

301 **DenseNet** is a generalization of ResNet that adds multiple residual connections to each
302 module. The beginning of a block of convolutions, is connected not only to the output
303 of that same module, but also to the outputs of all modules down-stream from it (Huang
304 et al., 2017).

305 **4.2.2 Attention-based Models**

306 This paper proposes using attention, rather than convolution, as the core model feature.
307 Attention is a deep learning mechanic that, rather than learn a weight per each input
308 pixel or a patch of pixels, learns a method for generating these weights from the input
309 data. Consequently, the models can weight each patch based on what its position is and
310 what the rest of the image depicts (Vaswani et al., 2017). In contrast, convolutions are
311 designed to analyze each patch of each input image using the same kernel of weights, re-
312 gardless of what the image depicts outside of the patch and what its position is. Formally,

313 convolutions enforce transition invariance, while attention models do not. Transition in-
314 variance in computer vision is achieved when the model maintains the same output even
315 if the objects in the image are moved around.

316 Attention’s ability to judge each image patch in the context of its position in the image
317 and the contents of the rest of the image is critical for making sound solar wind speed
318 predictions from the EUV data. First, the attention mechanism allows the model to as-
319 sign higher importance to features on the Sun’s surface if they appear in the equatorial
320 region. Moreover, the model is able to learn to distinguish between situations when an
321 active region interferes with a coronal hole, and when it does not. The weights it places
322 on the patches of the image with the coronal hole in it will depend not only on its po-
323 sition in the image, but also on whether the model identified an interference from an ac-
324 tive region. In contrast, convolution-based models were designed to identify an object
325 anywhere in the input image field. Therefore, they place equal weight on each image patch
326 as they process it using the same fixed-weight convolution kernel. It was assumed that
327 multiple layers of convolutions would learn increasingly complex representations by de-
328 riving higher-layer features from simple lower-layer ones. Recently, however, it was shown
329 that convolutional models do not recognize complex features, instead they aggregate low
330 level texture features from across the input image and then make their prediction based
331 on which texture prevails in the input image (Geirhos et al., 2018). Consequently, attention-
332 based models will make better and more theory-sensible predictions as it, for example,
333 will account for and internalize the higher importance of features in the equatorial re-
334 gion and the interference of active regions with coronal holes while convolution will fail
335 to do so.

336 The *Vision Transformer* was the first transformer architecture successfully used in im-
337 age recognition (Dosovitskiy et al., 2020). The architecture combines large image patches
338 with the attention mechanism. Each patch is first individually passed through a linear
339 projection, then the attention mechanism applies context-derived weights on each. The
340 result is then passed into two consecutive non-linear projections, sometimes called fully
341 connected layers, before being outputted. An important point of comparison is the size
342 of the model’s patches. While all benchmark models only consider patches of no more
343 than 5x5 pixels, our Vision Transformer works with patches of 16x16. This is meant to
344 allow it a larger receptive field and to steer clear of focusing on textures.

345 The *Transformer in Transformer* follows the same general architecture as the original
346 Vision Transformer, the crucial difference is that the linear projection at the beginning
347 of the outer transformer is replaced by an inner transformer that is modelled as a smaller

348 version of the same original Vision Transformer (Han et al., 2021). Therefore, the input
 349 image is first split into 16 by 16 patches. Each of these patches is then passed into the
 350 inner Vision Transformer, as if they were images in their own right. This splits them into
 351 smaller (4 x 4) patches still, derives the attention weight for each sub-patch based on the
 352 rest of each patch, and outputs the processed image back to the outer transformer. The
 353 outer transformer then uses these processed patches to derive its attention weights per
 354 each patch based on what the rest of the full image’s processed patches are like. Then
 355 the outer transformer uses two consecutive non-linear projections to produce the final
 356 output.

357 The *Swin Transformer* is similar to the Vision Transformer except it builds hierarchi-
 358 cal feature maps by merging image patches, as opposed to treating image patches sep-
 359 arately as in the Vision Transformer (Z. Liu et al., 2021). The idea is that the model is
 360 able to treat features on different scales, whereas the vanilla vision transformer is lim-
 361 iting itself to a predetermined patch size. Furthermore, a feature of the algorithmic con-
 362 struction is a linear scale in computational complexity based on image size.

363 These pre-trained attention-based models, as well as the benchmark CNN models, all
 364 accept three-channel RGB images normally. In order to use these powerful models, the
 365 solar images have to be repeated 3 times to form the three channels. Normally, one would
 366 use the advised normalisation schedule from the papers that produced these models. In
 367 this case however, since the models are not RGB in the first case, it was decided that
 368 an initial batch norm layer is applied before the model, so that the best normalisation
 369 schedule can be learned and not fixed. The reported RMSE and Correlation is then av-
 370 eraged over the five folds and reported.

$$371 \quad \text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - y_i)^2}$$

$$372 \quad \text{Correlation} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 (y_i - \bar{y})^2}}$$

373 where y_i is the real solar wind speed, x_i is the predicted solar wind speed, \bar{y}_i is the mean
 374 real speed, \bar{x}_i is the mean predicted speed, and n is the total number of data points.

375 **4.3 Other Experimental Details**

376 **4.3.1 Missing data**

377 Missing images are substituted with valid observations no more than 30 minutes removed
 378 from the missing datum. Missing solar wind speed data are interpolated from available
 379 data no more than 30 minutes removed. Time steps with no valid data for filling in the
 380 missing observations are discarded.

381 **4.3.2 Hyper-parameter selection**

382 Hyper-parameters are chosen using a Bayesian parameter sweep using the software Weights
383 & Biases (Biewald, 2020) based on the performance of the validation set. For cost rea-
384 sons, the sweep is conducted at 120 minutes resolution for only 30 epochs.

385 **4.3.3 Training process**

386 The loss function of the network is the default implementation of pytorch’s mean squared
387 error (squared L2 norm) (Paszke et al., 2019). The optimizer method to update the weights
388 of the network is the default implementation of the Adam optimizer in pytorch as well
389 (Kingma & Ba, 2014). Batch size is fixed at 64.

390 **4.3.4 Computation**

391 All experiments were run on V100 Nvidia GPU, resulting in a total compute of about
392 900 GPU hours.

393 **4.4 Year 2018 Evaluation**

394 Solar activity can vary significantly based on position in the solar cycle, so only testing
395 on 2018 only gives the performance of the model in that part of the solar cycle. It there-
396 fore cannot be representative of the generalisation of the model to other periods of the
397 solar cycle. However, Raju and Das (2021) provide results for a model trained on solar
398 imaging data with the entire year of 2018 held out for evaluation. As an extra exper-
399 iment and to compare to their study, a model will be trained with the training and test
400 set schedule shown in Figure 2b. Notably, Figure 2b features a 27 day test buffer before
401 the start of the 2018 test set. This buffer is present because of Raju and Das’ concern
402 of 27 day resurgence causing the training and test sets to not be independent. Our view
403 is that since this model is forecasting at a 4 day forecast, any image before that 4 days
404 could be used to train a model in a production system to make that 4 day forecast (es-
405 specially using the method of online learning). Despite the dependence, this 27-day old
406 image would be one of the most important images you would want to train on. Where
407 the dependence matters for forecasting purposes is crucially when the images are less than
408 the forecast horizon apart. This explains our choice of 4 day buffer otherwise. However,
409 for the point of comparison, this 27 day buffer is kept. Otherwise, all experimental pro-
410 cedures as detailed will remain the same as with the 5-fold split.

411 5 Results and Discussion

412 5.1 Comparison to Previous Works

413 Table 1 shows the comparison of our methodological and modelling pipeline, used with
 414 a range of feature extractors, against the most recent state of the art forecasting model
 415 in the field and two naive persistence model benchmarks. Notably, all of the models trained
 416 under our pipeline improve on the work by Upendran et al. (2020) by at least 8.8% in
 417 RMSE and 12.7% in correlation. Indeed, our pipeline with the GoogleNet feature ex-
 418 tractor, which is the same feature extractor as was used in the Upendran et al. (2020)
 419 model, demonstrated the total improvement our pipeline has delivered. It lowered the
 420 RMSE by 9.2% and increased the correlation by 14.6%. Furthermore, our best perform-
 421 ing model, based off the Swin Vision Transformer, improves on the state of the art by
 422 11.1% in RMSE and 17.4% in correlation. The model also outperforms at the 1, 2, and
 423 3 day time horizon because the 4 day forecast could also be used for those. Finally, trans-
 424 former feature extractors outperformed convolutional ones by about 1 to 2% in either
 425 metric when used in our model pipeline.

Table 1: Performance of our solar models compared to Upendran et al. (2020) forecasting solar wind speed using the EUV data at a 4 day forecast horizon in the period May 2010 to December 2018. Upendran et al. (2020).

Model	RMSE	% Improvement	Correlation	% Improvement
Persistence(4 day)	127.59	-57.1%	0.080	-85.2%
Persistence(27 day)	100.86	-24.2%	0.426	-21.1%
Upendran et al. (2020)	81.21	-	0.54	-
Our models				
Solar InceptionNet v4	74.09	8.8%	0.609	12.7%
Solar DenseNet	73.92	9.0%	0.611	13.1%
Solar GoogleNet	73.71	9.2%	0.619	14.6%
Solar ResNet	73.52	9.5%	0.618	14.4%
Solar TNT	72.70	10.5%	0.629	16.5%
Solar Vision Transformer	72.66	10.5%	0.630	16.7%
Solar Swin Transformer	72.21	11.1%	0.634	17.4%

426 Table 2 compares the performance of our best performing model, that is the one based
 427 on the Swin Transformer feature extractor, and the two persistence benchmarks against
 428 the predictions Raju and Das (2021) produced for the year 2018. This setup differs from
 429 that of table 1 in that table 1 tests the models on data examples sampled from the whole
 430 dataset, and thus across the solar cycle. The present comparison is made solely with re-
 431 spect to the solar cycle conditions present in the year 2018, as chosen by Raju and Das
 432 (2021). Our model shows a significant improvement of 8.3% in RMSE and 17.1% in cor-
 433 relation over the performance achieved by Raju and Das (2021).

Table 2: Performance of our solar models relative to (Raju & Das, 2021) predicting solar wind speed using EUV data at a 4 day forecast horizon in for the year 2018.

Model	RMSE	% Improvement	Correlation	% Improvement
Persistence(4 day)	118.76	-52.3%	-0.027	-104.9%
Persistence(27 day)	85.16	-9.2%	0.464	-15.6%
Raju and Das (2021)	78	-	0.55	-
Our model				
Solar Swin Transformer	71.65	8.3%	0.644	17.1%

434 5.2 Ablation Study

435 To demonstrate the stand-alone effect of our suggested techniques on the results, we con-
 436 ducted a study whereby each improvement is removed one at a time and the performance
 437 reduction reported. In the case of dropping the buffers, the no-buffer condition was im-
 438 plemented by making those buffers between the validation and training sets become part
 439 of the validation set, thus removing the separation between the two sets whilst adher-
 440 ing to a test-validation-train split that is comparable to that of the original condition.

441 Figure 4 shows that the dominant improvement has been the adjustment of the sampling
 442 frequency, excluding it causes 8.51% deterioration in RMSE and 9.70% in correlation.

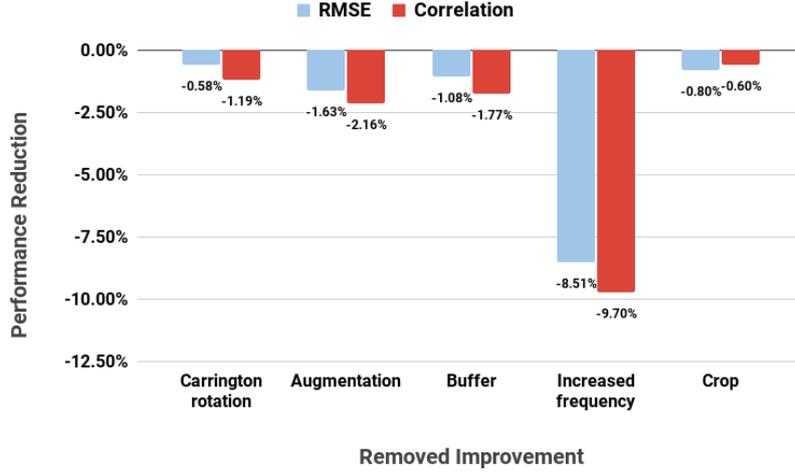


Figure 4: Performance reduction resulting from removing one improvement at a time.

443 Excluding the other 4 methodological, improvements delivers between 0.58% and 1.63%
 444 RMSE deterioration, and between 0.6% and 2.16% fall in correlation. While these fig-
 445 ures are modest in magnitude, it ought to be pointed out that the benefits appear un-
 446 correlated between the methods, and when they are all combined they deliver a signif-
 447 icant improvement over the previous works.

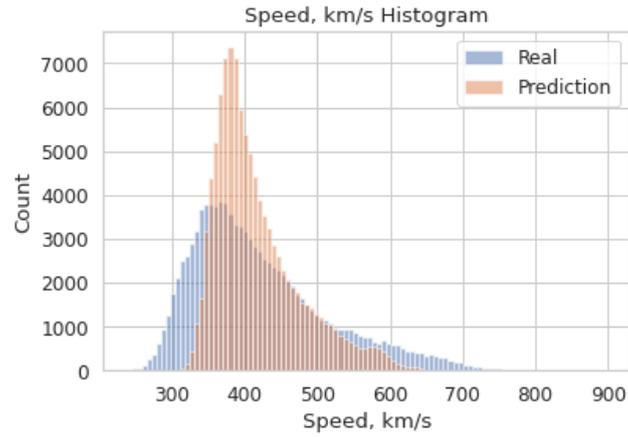
448 5.3 Prediction Analysis

449 Next, we analyze the predictions made by the best performing Swin Transformer model
 450 to get better understanding what aspects of the solar wind speed prediction task it gets
 451 right, and where, if at all, lie its systematic biases.

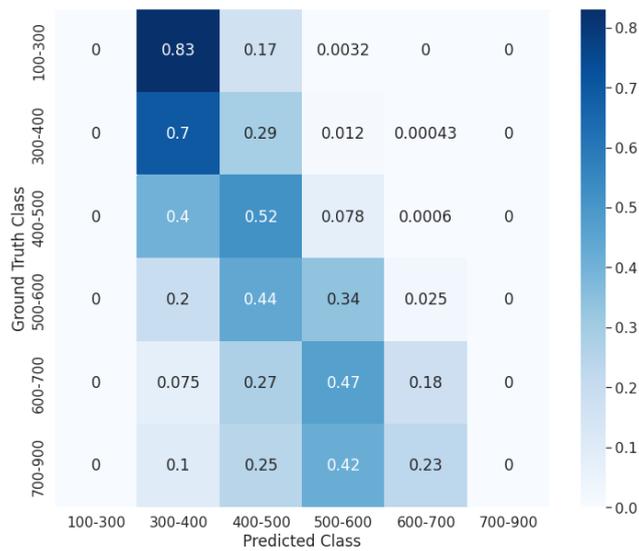
452 5.3.1 Distribution

453 Figure 5a shows the distributions of the solar wind speeds predicted by the top model
 454 and the underlying ground truth. Both distributions are roughly centered around the
 455 same mean with a positive skewness, i.e. they have long right-hand tails. The distribu-
 456 tions differ significantly in their kurtosis. The real data has lower kurtosis, that is it has
 457 more observations in both its right and left tails. The model’s predictions have notably
 458 higher kurtosis, as it has a much more pronounced peak at around its mean and much
 459 fewer observations in its tails. This is to be expected as the L2 loss function chosen, which
 460 all models in this domain use, is known to prioritize the average fit of the model over fit-
 461 ting the extremities. The distributions by themselves, however, do not tell the full story.
 462 For that we need to look at figure 5b, which shows the confusion matrix of binned speeds.
 463 Both predicted and actual solar wind speeds are split into 4 distinct class bins incremented

464 by 100 km/s and 2 catch all classes one at each extreme of the distributions. Each block
465 of the confusion matrix corresponds to one combination of a predicted class and a ground
466 truth, i.e. real, class. The value in the block represents the fraction of that real class that
467 were classified as the predicted class. Under a perfect prediction, the blocks would read
468 1.0 along the diagonal and 0 everywhere else. This would mean that all speeds were cor-
469 rectly predicted in their class. As it is however, our model shows a tendency to over-predict
470 the lower real speeds while under-predicting the higher speeds. Indeed, no speeds that
471 were in the 700-900 km/s range were correctly predicted as such. Similarly, no speeds
472 in the 100-300 range were correctly predicted. This confirms our suspicion that it is the
473 tail observations that are being regressed towards the mean that is driving both the er-
474 ror in the confusion matrix and the difference in the prediction and ground truth dis-
475 tributions.



(a) Distribution of predicted and real speeds



(b) Confusion matrix of binned speeds (km/s)

Figure 5: Distribution and confusion matrix of predicted speeds

476 **5.3.2 Coronal hole area**

477 It has been empirically established that there is a linear relationship between coronal hole
 478 area at low latitudes and peak solar wind speed (Nolte et al., 1976; Hofmeister et al., 2018).
 479 In order to test whether our model has learned this relationship we need to devise a way
 480 of obtaining images with specified coronal hole sizes at the desired latitude. We chose
 481 to generate our images using a background of enlarged uneventful solar region and a patch
 482 extracted from a coronal hole that can be sized as desired. Each patch size is moved across

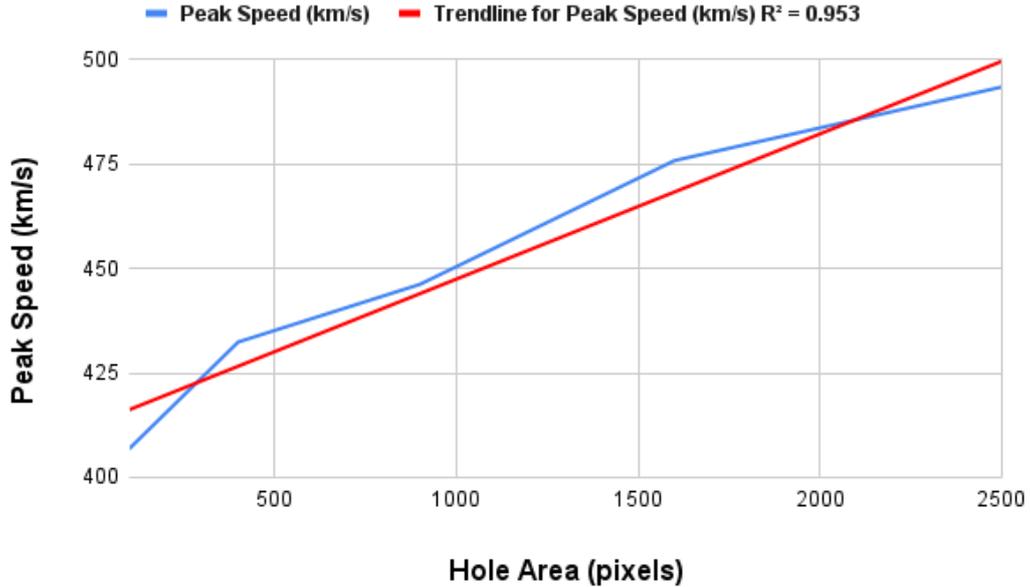


Figure 6: Peak speed of coronal holes at solar equator versus coronal hole area

483 the image, and the model's peak prediction for that size is recorded. Figure 6 plots the
 484 predicted solar wind speeds against the patch sizes. It shows that our model, indeed, suc-
 485 ceeded to learn the relationship established by Nolte et al. (1976) and Hofmeister et al.
 486 (2018) since its predictions follow very closely the empirically observed linear relation-
 487 ship with a high degree of correlation.

488 *5.3.3 Coronal hole intensity*

489 Obridko et al. (2009) found that the darker the coronal hole, the larger is the peak of
 490 the associated high speed stream. We test whether our model learned this empirical re-
 491 lationship by incrementally increasing the minimum brightness of a coronal hole. At each
 492 step, any pixel value below the minimum threshold is increased to the minimum value.
 493 Figure 7 shows the predicted speed for a large coronal hole visible on the day of 2016-
 494 12-06 at 00:00:00 am at various minimum intensities. As we increase the brightness of
 495 the coronal hole, the model starts to forecast lower solar wind speeds. This suggests that
 496 the model has learned the Obridko et al. (2009) empirical relationship that the darker
 497 the hole, the stronger the solar wind.

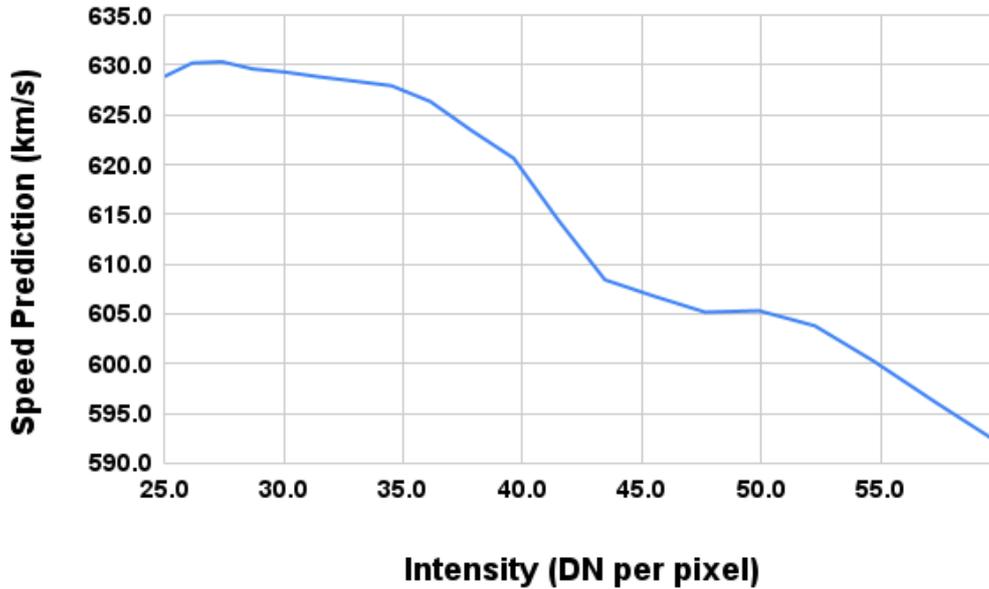
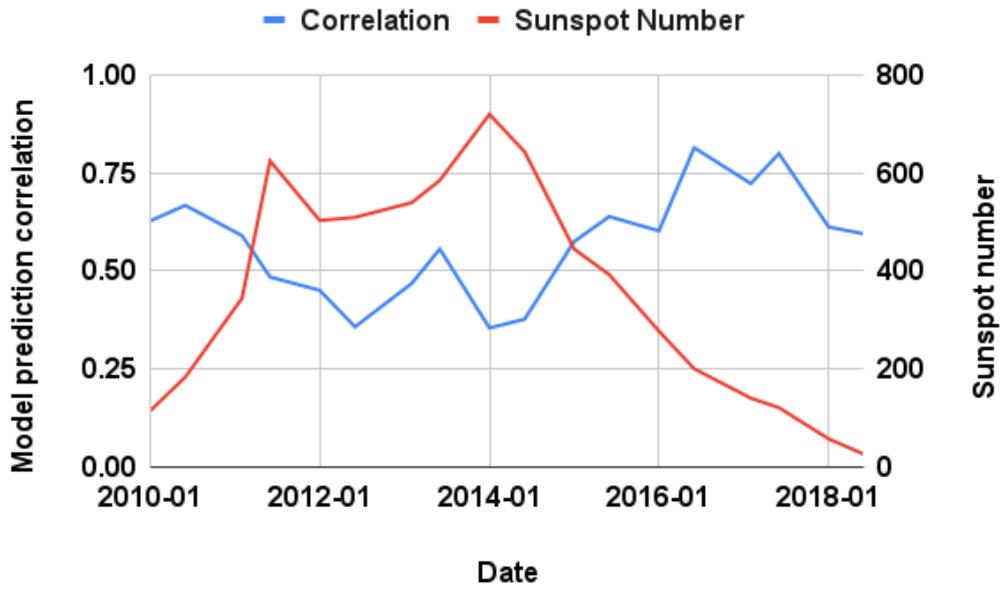


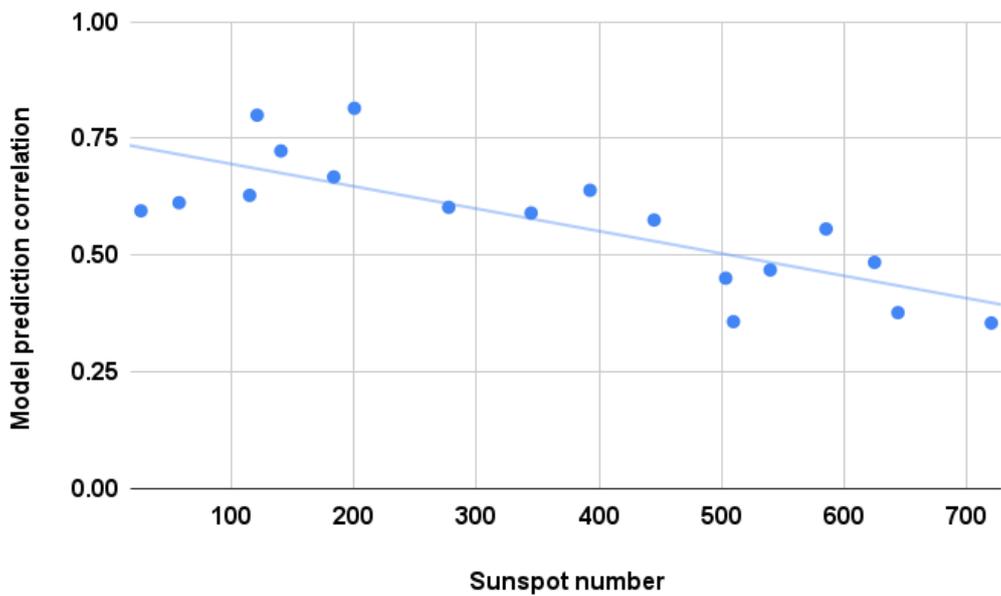
Figure 7: Speed prediction vs minimum pixel intensity for a coronal hole pictured on 2016-12-10

498 *5.3.4 Solar cycle variability*

499 The performance of the model is highly dependent on position in the solar cycle. Fig-
 500 ure 8a plots the correlation of the model prediction with the ground truth (blue) at 6
 501 month intervals against the number of sunspots (red) in that period. The sunspot num-
 502 ber represents the solar cycle. Notably, the model's prediction correlation to ground truth
 503 is much better in the declining phase of the solar cycle, that is in the 2016 to 2018. At
 504 the same time, it performs much worse around the peak of the solar cycle in 2014. This
 505 relationship is confirmed when we view the data as correlation-sunspot number couples
 506 and visualize them in a scatter plot. This is shown in the figure 8b. We observe a strong,
 507 0.78, negative correlation of the number of sunspots and the model prediction correla-
 508 tion to the ground truth. Since sunspot number is used to measure the solar cycle, this
 509 suggests that the model performance is highly dependent on the solar cycle and more
 510 specifically on the prevalent type of solar activity in a given period.



(a) Model prediction correlation and sunspot number vs date



(b) Model prediction correlation and sunspot number

Figure 8: Model performance compared to sunspot number

511 Indeed, a key component of the model's performance across the solar cycle is the type
 512 of encountered solar features. The top two panels of the figure 9 show the model's per-
 513 formance in early 2012, with 80.81 RMSE and 0.45 correlation, and in late 2016, with
 514 73.32 RMSE and 0.81 correlation. The solar wind behaviour in the later half of 2016,
 515 was driven by coronal holes and the high speed solar wind streams associated with them.
 516 Whereas, 2012 had a much higher sunspot number and had far more Earth-directed CMEs.

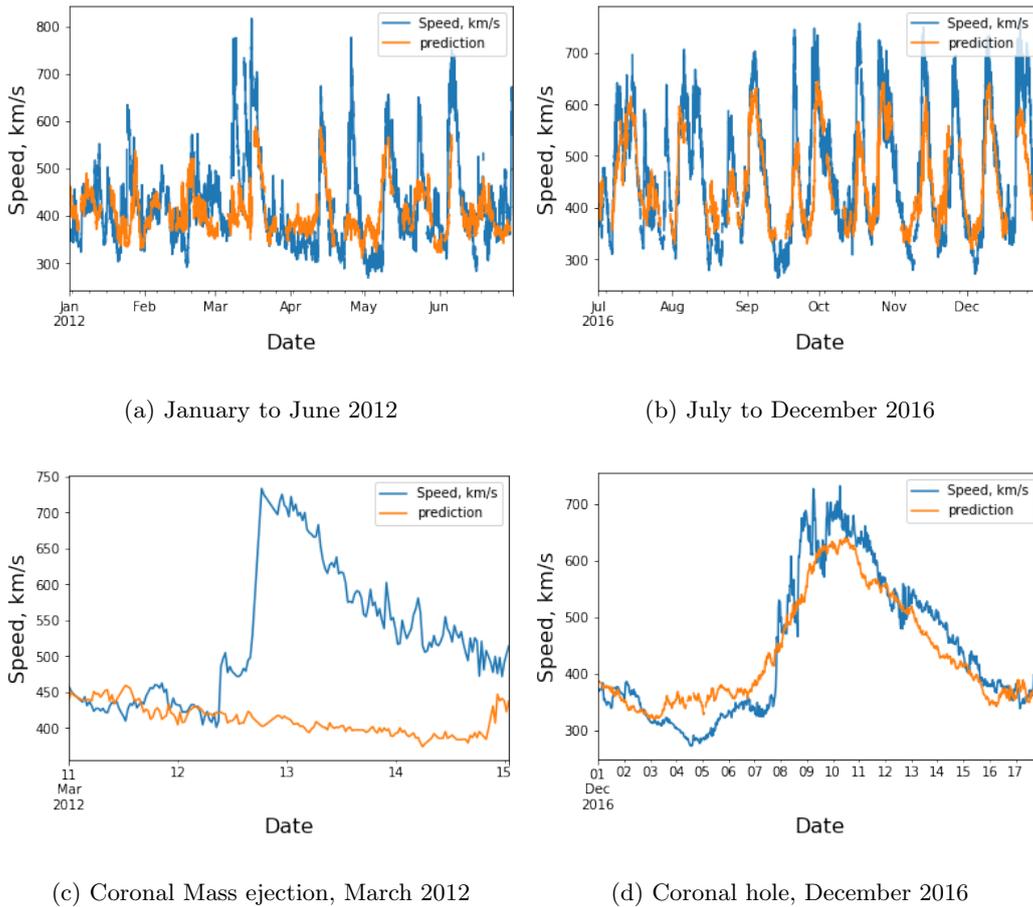


Figure 9: Solar Swin Transformer performance in different parts of the solar cycle and on different solar phenomena

517 We observe a marked difference in performance between predictions driven by CMEs and
 518 coronal holes. Figures 9d and 9c show how the model captures the longer lasting, speed
 519 profile of a coronal hole quite well, while missing the speed profile of the sudden CME.
 520 This offers an explanation to the pronounced variability in the model's prediction qual-
 521 ity. The solar activity in the declining phase is driven by coronal holes. These are more
 522 easily picked up by the models. Since the Sun in the later half of 2016 was in the declin-

523 ing phase, the models' performance was much better. In 2012, a year with far more CMEs,
524 the model performance was reduced, as the models struggled to catch the CMEs.

525 The failure to fit on the more sudden coronal mass ejections is a chief limitation of the
526 models developed in this space. It can be ascribed to the lack of significant and persis-
527 tent CME-related features in the EUV images, preventing them from being captured by
528 the models. We note that ML models using solar EUV images alone to forecast other
529 space weather related parameters such as geomagnetic activity as measured by the AE
530 or Kp indices or suprathermal electrons at geostationary orbit would most likely suffer
531 from the same limitation resulting in a similar pattern of behaviour with the best cor-
532 relations during the declining phase of the solar cycle and the worst correlations around
533 solar maximum.

534 **6 Conclusions**

535 This study uses attention-based machine vision models and a set of methodological and
536 modelling improvements to forecast the solar wind speed at L1 using solar images at 211Å
537 wavelength. These improvements result in 11.1% lower RMSE and 17.4% higher predic-
538 tion correlation with the ground truth when compared to previous works. Additionally,
539 this study observed that attention-based architectures in general have about 2-3% per-
540 formance edge in both RMSE and correlation over the previously-used convolutional al-
541 ternatives. The model's performance is highly dependent on the position in the solar cy-
542 cle. The model performance is strongly negatively correlated with the sunspot number,
543 as the model performance is better in the declining phase of the solar cycle when the so-
544 lar wind behaviour is dominated by coronal hole activity. Finally, the model has inde-
545 pendently learned two empirical relationships between coronal established by previous
546 publications. First, it complies with the observed linear relationship between coronal hole
547 area and the peak speed associated with it. Second, it learned that the darker the coro-
548 nal hole, the stronger the solar wind speed associated with it.

549 **Open Research**

550 The SDOML 211 Å image data is available here: <https://purl.stanford.edu/vk217bh4910>

551 The OmniWeb solar wind data is available here: https://omniweb.gsfc.nasa.gov/form/omni_min.html

552
553 Code for analysing model output is supplied here: [https://github.com/eddbrown/solar](https://github.com/eddbrown/solar-swin-transformer-output-data)
554 `-swin-transformer-output-data`

555 **Acknowledgements**

556 For this study, we acknowledge extensive use of the SDOML (Galvez et al., 2019) dataset.

557 These images are processed versions of images taken by the AIA (Atmospheric Imaging
558 Assembly) instrument aboard the Solar Dynamic Observatory.

559 Furthermore we acknowledge the use of the OMNIWeb service at their provided mea-
560 surements of the solar wind speed at L1.

561 Regarding software, we acknowledge extensive use of python and the python packages
562 numpy (Harris et al., 2020) and pytorch (Paszke et al., 2019).

563 For experiment tracking and general project organisation, we acknowledge extensive use
564 of the software provided by Weights and Biases (Biewald, 2020).

565 N. P. Meredith and R. B. Horne would like to acknowledge support from the NERC grants
566 NE/V00249X/1 (Sat-Risk) and NE/R016038/1.

567 **References**

- 568 Bailey, R. L., Reiss, M. A., Arge, C. N., Möstl, C., Henney, C. J., Owens, M. J.,
 569 ... Hinterreiter, J. (2021). Using gradient boosting regression to im-
 570 prove ambient solar wind model predictions. *Space Weather*, 19(5),
 571 e2020SW002673. (e2020SW002673 2020SW002673) doi: [https://doi.org/](https://doi.org/10.1029/2020SW002673)
 572 10.1029/2020SW002673
- 573 Biewald, L. (2020). *Experiment tracking with weights and biases*. (Software available
 574 from wandb.com)
- 575 Borovsky, J. E., & Denton, M. H. (2006). Differences between cme-driven storms
 576 and cir-driven storms. *Journal of Geophysical Research: Space Physics*,
 577 111(A7). doi: <https://doi.org/10.1029/2005JA011447>
- 578 Burlaga, L., & Lepping, R. (1977). The causes of recurrent geomagnetic storms.
 579 *Planetary and Space Science*, 25(12), 1151-1160. doi: [https://doi.org/10.1016/](https://doi.org/10.1016/0032-0633(77)90090-3)
 580 0032-0633(77)90090-3
- 581 Chandorkar, M., Furtlehner, C., Poduval, B., Camporeale, E., & Sebag, M. (2019).
 582 Dynamic time lag regression: predicting what & when. In *International confer-*
 583 *ence on learning representations*.
- 584 Chen, G.-m., Xu, J., Wang, W., Lei, J., & Burns, A. G. (2012). A comparison of the
 585 effects of cir- and cme-induced geomagnetic activity on thermospheric densities
 586 and spacecraft orbits: Case studies. *Journal of Geophysical Research: Space*
 587 *Physics*, 117(A8). doi: <https://doi.org/10.1029/2012JA017782>
- 588 Cranmer, S. R. (2009). Coronal holes. *Living Reviews in Solar Physics*, 6, 3,66. doi:
 589 <https://doi.org/10.12942/lrsp-2009-3>
- 590 Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet:
 591 A large-scale hierarchical image database. In *2009 ieee conference on computer*
 592 *vision and pattern recognition* (p. 248-255). doi: 10.1109/CVPR.2009.5206848
- 593 Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner,
 594 T., ... Houlsby, N. (2020). An image is worth 16x16 words: Transformers for
 595 image recognition at scale. *CoRR*, abs/2010.11929.
- 596 Galvez, R., Fouhey, D. F., Jin, M., Szenicer, A., Muñoz-Jaramillo, A., Cheung,
 597 M. C. M., ... Thomas, R. (2019, may). A machine-learning data set prepared
 598 from the NASA solar dynamics observatory mission. *The Astrophysical Journal*
 599 *Supplement Series*, 242(1), 7. doi: 10.3847/1538-4365/ab1005
- 600 Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., & Brendel,
 601 W. (2018). Imagenet-trained cnns are biased towards texture; increasing shape
 602 bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*.
- 603 Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., & Wang, Y. (2021). *Transformer in*

- 604 *transformer*.
- 605 Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cour-
 606 napeau, D., ... Oliphant, T. E. (2020, September). Array programming with
 607 NumPy. *Nature*, 585(7825), 357–362. doi: 10.1038/s41586-020-2649-2
- 608 He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image
 609 recognition. In *2016 IEEE conference on computer vision and pattern recognition*
 610 (*cvpr*) (p. 770-778). doi: 10.1109/CVPR.2016.90
- 611 Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural compu-*
 612 *tation*, 9(8), 1735–1780.
- 613 Hofmeister, S. J., Veronig, A. M., Temmer, M., Vennerstrom, S., Heber, B., &
 614 Vrsnak, B. (2018). The dependence of the peak velocity of high-speed so-
 615 lar wind streams as measured in the ecliptic by ace and the stereo satellites
 616 on the area and co-latitude of their solar source coronal holes. *Journal of*
 617 *Geophysical Research. Space Physics*, 123, 1738 - 1753.
- 618 Horne, R. B., Phillips, M. W., Glauert, S. A., Meredith, N. P., Hands, A. D. P., Ry-
 619 den, K. A., & Li, W. (2018). Realistic worst case for a severe space weather
 620 event driven by a fast solar wind stream. *Space Weather*, 16(9), 1202-1215.
 621 doi: <https://doi.org/10.1029/2018SW001948>
- 622 Huang, G., Liu, Z., & Weinberger, K. Q. (2017). Densely connected convolutional
 623 networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition*
 624 (*CVPR*), 2261-2269.
- 625 Kilpua, E., Balogh, A., Von Steiger, R., & Liu, Y. (2017, 01). Geoeffective proper-
 626 ties of solar transients and stream interaction regions. In (p. 295-338). doi: 10
 627 .1007/978-94-024-1588-9_9
- 628 Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*
 629 *preprint arXiv:1412.6980*.
- 630 Kozyra, J. U., Crowley, G., Emery, B. A., Fang, X., Maris, G., Mlynczak, M. G., ...
 631 Yee, J.-H. (2006). Response of the upper/middle atmosphere to coronal holes
 632 and powerful high-speed solar wind streams in 2003. In *Recurrent magnetic*
 633 *storms: Corotating solar wind streams* (p. 319-340). American Geophysical
 634 Union (AGU). doi: <https://doi.org/10.1029/167GM24>
- 635 Lemen, J. R., Akin, D. J., Boerner, P. F., Chou, C., Drake, J. F., Duncan, D. W.,
 636 ... others (2011). The atmospheric imaging assembly (aia) on the solar dy-
 637 namics observatory (sdo). In *The solar dynamics observatory* (pp. 17–40).
 638 Springer.
- 639 Liu, D. D., Huang, C., Lu, J. Y., & Wang, J. S. (2011, 05). The hourly average solar
 640 wind velocity prediction based on support vector regression method. *Monthly*

- 641 *Notices of the Royal Astronomical Society*, 413(4), 2877-2882. doi: 10.1111/j
642 .1365-2966.2011.18359.x
- 643 Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., ... Guo, B. (2021). Swin
644 transformer: Hierarchical vision transformer using shifted windows. *CoRR*,
645 *abs/2103.14030*.
- 646 Meredith, N. P., Horne, R. B., Lam, M. M., Denton, M. H., Borovsky, J. E., &
647 Green, J. C. (2011). Energetic electron precipitation during high-speed solar
648 wind stream driven storms. *Journal of Geophysical Research: Space Physics*,
649 116(A5). doi: <https://doi.org/10.1029/2010JA016293>
- 650 Nolte, J., Krieger, A., Timothy, A., Gold, R., Roelof, E., Vaiana, G., ... McIntosh,
651 P. (1976). Coronal holes as sources of solar wind. *Solar Physics*, 46(2),
652 303–322.
- 653 Obridko, V., Shelting, B., Livshits, I., & Askerov, A. (2009). Relationship between
654 the contrast of coronal holes and parameters of the solar wind streams. *Astron-
655 omy reports*, 53(11), 1050–1058.
- 656 Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., ... Chintala,
657 S. (2019). Pytorch: An imperative style, high-performance deep learning li-
658 brary. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, &
659 R. Garnett (Eds.), *Advances in neural information processing systems 32* (pp.
660 8024–8035). Curran Associates, Inc.
- 661 Raju, H., & Das, S. (2021). Cnn-based deep learning model for solar wind forecast-
662 ing. *Sol Phys*. doi: <https://doi.org/10.1007/s11207-021-01874-6>
- 663 Richardson, I. G., Webb, D. F., Zhang, J., Berdichevsky, D. B., Biesecker, D. A.,
664 Kasper, J. C., ... Zhukov, A. N. (2006). Major geomagnetic storms
665 (dst less than or equal to -100 nt) generated by corotating interaction re-
666 gions. *Journal of Geophysical Research: Space Physics*, 111(A7). doi:
667 <https://doi.org/10.1029/2005JA011476>
- 668 Ridpath, I. (2012). *A dictionary of astronomy*. Oxford University Press.
- 669 St. Cyr, O. C., Howard, R. A., Sheeley Jr., N. R., Plunkett, S. P., Michels, D. J.,
670 Paswaters, S. E., ... Lamy, P. L. (2000). Properties of coronal mass
671 ejections: Soho lasco observations from january 1996 to june 1998. *Jour-
672 nal of Geophysical Research: Space Physics*, 105(A8), 18169-18185. doi:
673 <https://doi.org/10.1029/1999JA000381>
- 674 Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S. E., Anguelov, D., ... Rabi-
675 novich, A. (2014). Going deeper with convolutions. *CoRR*, *abs/1409.4842*.
- 676 Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethink-
677 ing the inception architecture for computer vision. *2016 IEEE Conference on*

- 678 *Computer Vision and Pattern Recognition (CVPR)*, 2818-2826.
- 679 Tsurutani, B. T., Gonzalez, W. D., Gonzalez, A. L. C., Tang, F., Arballo, J. K., &
680 Okada, M. (1995). Interplanetary origin of geomagnetic activity in the declin-
681 ing phase of the solar cycle. *Journal of Geophysical Research: Space Physics*,
682 *100*(A11), 21717-21733. doi: <https://doi.org/10.1029/95JA01476>
- 683 Turner, N. E., Mitchell, E. J., Knipp, D. J., & Emery, B. A. (2006). En-
684 ergetics of magnetic storms driven by corotating interaction regions: A
685 study of geoeffectiveness. In *Recurrent magnetic storms: Corotating so-*
686 *lar wind streams* (p. 113-124). American Geophysical Union (AGU). doi:
687 <https://doi.org/10.1029/167GM11>
- 688 Upendran, V., Cheung, M. C. M., Hanasoge, S., & Krishnamurthi, G. (2020). Solar
689 wind prediction using deep learning. *Space Weather*, *18*(9), e2020SW002478.
690 (e2020SW002478 10.1029/2020SW002478) doi: [https://doi.org/10.1029/](https://doi.org/10.1029/2020SW002478)
691 [2020SW002478](https://doi.org/10.1029/2020SW002478)
- 692 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ...
693 Polosukhin, I. (2017). Attention is all you need..
- 694 Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind
695 velocity from solar magnetic field observations using hybrid intelligent
696 systems. *Physics and Chemistry of the Earth*, *22*(7), 617-622. (Low-
697 Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi:
698 [https://doi.org/10.1016/S0079-1946\(97\)00186-9](https://doi.org/10.1016/S0079-1946(97)00186-9)
- 699 Wintoft, P., & Lundstedt, H. (1999). A neural network study of the mapping from
700 solar magnetic fields to the daily average solar wind velocity. *Journal of Geo-*
701 *physical Research: Space Physics*, *104*(A4), 6729-6736. doi: [https://doi.org/10](https://doi.org/10.1029/1998JA900183)
702 [.1029/1998JA900183](https://doi.org/10.1029/1998JA900183)
- 703 Yang, Y., Shen, F., Yang, Z., & Feng, X. (2018). Prediction of solar wind speed at
704 1 au using an artificial neural network. *Space Weather*, *16*(9), 1227-1244. doi:
705 <https://doi.org/10.1029/2018SW001955>