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

Edward E. J. Brown¹, Filip Svoboda¹, Nigel P. Meredith², Nicholas Lane¹, and Richard B. Horne²

¹University of Cambridge ²British Antarctic Survey

November 22, 2022

Abstract

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

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

Edward J. E. Brown ^{1,2,*}, Filip Svoboda ^{1,*}, Nigel P. Meredith², Nicholas Lane ^{1,3}, Richard B. Horne²

¹ Department of Computer Science and Technology, University of Cambridge, Cambridge, UK
$^2\mathrm{British}$ Antarctic Survey, Natural Environment Research Council, Cambridge, UK
³ Samsung AI Center, Cambridge, UK
*equal contribution

4

5

Key Points: Attention-based machine vision models and methodological enhancements are developed to improve solar wind speed forecasts from solar images Attention-based architectures outperform convolutional models, motivating their use in future studies and production systems The models perform best in the declining phase of the solar cycle when activity 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
- solar wind speed a difficult, high-impact, and unsolved problem. At a four day time
- horizon, this study uses attention-based models and a set of methodological improvements
- to deliver an 11.1% lower RMSE error and a 17.4% higher prediction correlation com-
- pared to the previous work testing on the period from 2010 to 2018. Our analysis shows
- that attention-based models combined with our pipeline consistently outperform con-
- 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
- the solar cycle, with the best performance occurring in the declining phase.

²⁹ Plain language summary

³⁰ Solar images contain rich information that can be used to forecast conditions at Earth.

³¹ This study develops a robust methodology for processing solar images and trains ma-

s2 chine 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 qui-

 $_{35}$ eter, declining phase of the solar cycle when the solar activity is driven by coronal holes.

 $_{36}$ Finally, the trained models learn properties of coronal holes that agree with prior em-

³⁷ pirical studies.

38 1 Introduction

The solar wind is a stream of charged particles that is emitted from the upper atmosphere 39 of the Sun. The speed, density, temperature and the magnitude and direction of the as-40 sociated magnetic field of the solar wind are constantly varying affecting the way in which 41 it ultimately interacts with the Earth's magnetosphere. High speed solar wind streams 42 (HSS) emanating from coronal holes are particularly effective at coupling with the Earth's 43 magnetosphere. The weak storms they produce tend to have long-lasting recovery phases 44 which often result in prolonged and enhanced substorm activity (Tsurutani et al., 1995; 45 Meredith et al., 2011). This results in repeated injections of suprathermal electrons into 46 the inner magnetosphere and significant increases in the fluxes of relativistic electrons 47 in the outer radiation belt, increasing the risk to satellites via surface charging and in-48 ternal charging respectively (e.g., Borovsky and Denton (2006)). Indeed, it has been sug-49 gested that satellites at geostationary orbit are more likely to be at risk from an extreme 50 HSS-driven storm than a Carrington type event (Horne et al., 2018). Furthermore, pro-51 longed and enhanced substorm activity associated with HSS-driven storms results in in-52 creased thermospheric densities and satellite drag (Chen et al., 2012). Consequently, ac-53 curately forecasting the solar wind speed associated with coronal holes is very impor-54 tant for our modern society. 55

Coronal holes are large dark areas on the Sun as seen in extreme ultraviolet (EUV) and 56 soft X-ray images (Cranmer, 2009). They are regions of open magnetic field and cooler 57 plasma, leading to the production of high speed solar wind streams. Coronal holes are 58 long-lasting features that can persist from one solar rotation to the next, giving rise to 59 a 27 day periodicity in the arrival of HSS at Earth. The occurrence rate of coronal holes 60 peaks during the declining phase of the solar cycle (Burlaga & Lepping, 1977) and high 61 speed streams observed at Earth during these intervals tend to be coronal-hole driven. 62 The distribution of speeds in high speed streams associated with coronal holes ranges 63 from 400 to 800 kms^{-1} (Kilpua et al., 2017). While these streams do not result in ma-64

-4-

jor geomagnetic storms (Richardson et al., 2006), they have extensive recovery phases,

typically lasting from 5–10 days, and, as a result, may deposit more energy in the mag-

netosphere than larger storms (Kozyra et al., 2006; Turner et al., 2006).

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

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

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

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

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

⁹⁹ 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.

¹⁰⁶ 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-imagebased 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 so-114 lar EUV images using deep learning techniques. The work uses images from both 193 115 and 211 Angstrom wavelengths to forecast the solar wind speed at a one day resolution. 116 Upendran uses GoogleNet (Szegedy et al., 2014), trained on the ImageNet dataset (Deng 117 et al., 2009), as a feature extractor for each image. The extracted per-image features are 118 then passed into an LSTM Recurrent Neural Network (Hochreiter & Schmidhuber, 1997) 119 to produce the predicted solar wind speed. The study achieves a best performing model 120 at a lag of 3 days and a history of 4 days, with a correlation of 0.55 and an RMSE 80.28 121 km/s. This study will build on this insightful initial work. 122

Next, Raju and Das (2021) proposed a smaller three-layer convolutional feature extrac-123 tor, which they train on the 193 Angstrom wavelength solar EUV images. Their method 124 targets a subtly different task than that of Upendran et al. (2020). While Upendran et 125 al. (2020) forecast future solar wind speeds based on images at a fixed distance in the 126 present, Raju and Das (2021) backcast current solar wind speed based on flexible-lag past 127 images. Specifically, Raju and Das (2021) use the current solar wind speed to infer which 128 past image was likely to have caused the recorded solar wind speed, and then pass this 129 image into their model with the expectation that the model will be able to correctly re-130 construct the observed solar wind speed. The key difference between the two approaches 131 is that in the forecasting setup the model needs to use the image information to both 132 infer the speed of the caused solar wind and make a judgement on whether the parti-133 cle stream at the given speed will be geo-effective. In the Raju and Das (2021) backcast-134

-6-

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

155 **3 Data**

156 3.1 Solar Images

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

¹⁶⁵ 3.2 Solar Wind Speed

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

-7-

- ¹⁷⁰ The solar wind speed is highly auto-correlated with itself over hourly time periods and
- 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
- a synodic rotation period of approximately 27 days and some longer lasting features, such
- 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
- in training and test sets due to their dependence on each-other. This is further discussed
- in Section 4.1.7.

178 4 Methodology

179 4.1 Methodological Improvements

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

182 4.1.1 Image pre-processing

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

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

198 4.1.2 Sampling frequency

- We replace the previously used daily sampling resolution with a 30 minute schedule, because 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-
- 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)

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

208 4.1.4 North-south augmentation

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

213 4.1.5 Single image versus sequence

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

224 4.1.6 Feature extractor re-training

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

230 4.1.7 Training, validation and test sets

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



(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

4.2 Model Architectures

For this study, the architectures for the different models will follow the format in Fig-

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
- this model as well as the solar wind speed from one Carrington rotation ago are then passed
- 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
- 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

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

Previous work used convolutional models in the forecasting of solar wind, (Upendran et al. (2020); Raju and Das (2021)). These models are designed to process images, each of which has three dimensions - the height, the width, and the number of channels. A standard colour image has 3 channels: red, green, and blue. Convolutions are operations that split the image into a grid of patches and then use a three dimensional kernel to compute weighted averages per each patch. The same kernel is used on each patch and the averages it produces become the pixel values the layer outputs. Multiple kernels may be employed, in which case their outputs are treated as separate channels of the outputted image.

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

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

ResNet is a predecessor of GoogleNet. ResNet's modules consist of two consecutive convolutions, and a so called residual connection. The residual connection is a bypass that circumvents the two convolutions. In effect this results in a block that outputs both its convolution's output as well as the original inputs to the block. This trick helps to propagate the training gradients through the network, mitigating the vanishing gradient problem. The architecture was the first one to breach the 20 layer depth ceiling (He et al., 2016).

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

305 4.2.2 Attention-based Models

This paper proposes using attention, rather than convolution, as the core model feature. Attention is a deep learning mechanic that, rather than learn a weight per each input pixel or a patch of pixels, learns a method for generating these weights from the input data. Consequently, the models can weight each patch based on what its position is and what the rest of the image depicts (Vaswani et al., 2017). In contrast, convolutions are designed to analyze each patch of each input image using the same kernel of weights, regardless of what the image depicts outside of the patch and what its position is. Formally, convolutions enforce transition invariance, while attention models do not. Transition invariance in computer vision is achieved when the model maintains the same output even if the objects in the image are moved around.

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

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

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

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

The Swin Transformer is similar to the Vision Transformer except it builds hierarchical feature maps by merging image patches, as opposed to treating image patches separately as in the Vision Transformer (Z. Liu et al., 2021). The idea is that the model is able to treat features on different scales, whereas the vanilla vision transformer is limiting itself to a predetermined patch size. Furthermore, a feature of the algorithmic construction is a linear scale in computational complexity based on image size.

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

RMSE =
$$\sqrt{\frac{1}{n} \sum_{i=1}^{n} \left(x_i - y_i\right)^2}$$

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

where y_i is the real solar wind speed, x_i is the predicted solar wind speed, $\overline{y_i}$ is the mean real speed, $\overline{x_i}$ is the mean predicted speed, and n is the total number of data points.

4.3 Other Experimental Details

376 4.3.1 Missing data

3

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

381 4.3.2 Hyper-parameter selection

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

385 4.3.3 Training process

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

390 4.3.4 Computation

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

³⁹³ 4.4 Year 2018 Evaluation

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

411 5 Results and Discussion

5.1 Comparison to Previous Works

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

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

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

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
- (2021). Our model shows a significant improvement of 8.3% in RMSE and 17.1% in cor-

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

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

Figure 4 shows that the dominant improvement has been the adjustment of the sampling 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.

Excluding the other 4 methodological, improvements delivers between 0.58% and 1.63% RMSE deterioration, and between 0.6% and 2.16% fall in correlation. While these figures are modest in magnitude, it ought to be pointed out that the benefits appear uncorrelated between the methods, and when they are all combined they deliver a significant improvement over the previous works.

448 5.3 Prediction Analysis

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

452 5.3.1 Distribution

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

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

475

-20-



(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

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

- of obtaining images with specified coronal hole sizes at the desired latitude. We chose
- to generate our images using a background of enlarged uneventful solar region and a patch
- 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

the image, and the model's peak prediction for that size is recorded. Figure 6 plots the predicted solar wind speeds against the patch sizes. It shows that our model, indeed, succeeded to learn the relationship established by Nolte et al. (1976) and Hofmeister et al. (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

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



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

498 5.3.4 Solar cycle variability

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



(a) Model prediction correlation and sunspot number vs date



(b) Model prediction correlation and sunspot number

Figure 8: Model performance compared to sunspot number

Indeed, a key component of the model's performance across the solar cycle is the type

of encountered solar features. The top two panels of the figure 9 show the model's per-

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,

- 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.



(c) Coronal Mass ejection, March 2012 (d) Coronal hole, December 2016

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

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

ing phase, the models' performance was much better. In 2012, a year with far more CMEs, 523

the model performance was reduced, as the models struggled to catch the CMEs. 524

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

6 Conclusions 534

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

Open Research 549

550

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

The OmniWeb solar wind data is available here: https://omniweb.gsfc.nasa.gov/form/ 551 omni_min.html 552

Code for analysing model output is supplied here: https://github.com/eddbrown/solar 553 -swin-transformer-output-data 554

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
- 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
- numpy (Harris et al., 2020) and pytorch (Paszke et al., 2019).
- ⁵⁶³ For experiment tracking and general project organisation, we acknowledge extensive use
- 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.

References

567

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/
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/
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, Gm., 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, LJ., 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., Houls by, 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, JH. (2006). Response of the upper/middle atmosphere to coronal holes
632	and powerful high-speed solar wind streams in 2003. In <i>Recurrent magnetic</i>
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	\dots others (2011). The atmospheric imaging assembly (aia) on the solar dy-
637	namics observatory (sdo). In <i>The solar dynamics observatory</i> (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	$(e2020SW002478\ 10.1029/2020SW002478) \qquad doi:\ https://doi.org/10.1029/$
691	2020SW002478
692	Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N.,
692 693	Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need
692 693 694	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind
692 693 694 695	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent
692 693 694 695 696	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low-
692 693 694 695 696 697	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low-Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi:
692 693 694 695 696 697 698	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low-Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi: https://doi.org/10.1016/S0079-1946(97)00186-9
692 693 694 695 696 697 698 699	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low-Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi: https://doi.org/10.1016/S0079-1946(97)00186-9 Wintoft, P., & Lundstedt, H. (1999). A neural network study of the mapping from
 692 693 694 695 696 697 698 699 700 	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low- Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi: https://doi.org/10.1016/S0079-1946(97)00186-9 Wintoft, P., & Lundstedt, H. (1999). A neural network study of the mapping from solar magnetic fields to the daily average solar wind velocity. <i>Journal of Geo</i>-
692 693 694 695 696 697 698 699 700	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low- Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi: https://doi.org/10.1016/S0079-1946(97)00186-9 Wintoft, P., & Lundstedt, H. (1999). A neural network study of the mapping from solar magnetic fields to the daily average solar wind velocity. <i>Journal of Geo- physical Research: Space Physics</i>, 104(A4), 6729-6736. doi: https://doi.org/10
 692 693 694 695 696 697 698 699 700 701 702 	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low- Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi: https://doi.org/10.1016/S0079-1946(97)00186-9 Wintoft, P., & Lundstedt, H. (1999). A neural network study of the mapping from solar magnetic fields to the daily average solar wind velocity. <i>Journal of Geo- physical Research: Space Physics</i>, 104(A4), 6729-6736. doi: https://doi.org/10 .1029/1998JA900183
 692 693 694 695 696 697 698 699 700 701 702 703 	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low- Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi: https://doi.org/10.1016/S0079-1946(97)00186-9 Wintoft, P., & Lundstedt, H. (1999). A neural network study of the mapping from solar magnetic fields to the daily average solar wind velocity. <i>Journal of Geo- physical Research: Space Physics</i>, 104(A4), 6729-6736. doi: https://doi.org/10 .1029/1998JA900183 Yang, Y., Shen, F., Yang, Z., & Feng, X. (2018). Prediction of solar wind speed at
 692 693 694 695 696 697 698 699 700 701 702 703 704 	 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention is all you need Wintoft, P., & Lundstedt, H. (1997). Prediction of daily average solar wind velocity from solar magnetic field observations using hybrid intelligent systems. <i>Physics and Chemistry of the Earth</i>, 22(7), 617-622. (Low- Altitude Investigation of Dayside Magnetospheric Boundary Processes) doi: https://doi.org/10.1016/S0079-1946(97)00186-9 Wintoft, P., & Lundstedt, H. (1999). A neural network study of the mapping from solar magnetic fields to the daily average solar wind velocity. <i>Journal of Geo- physical Research: Space Physics</i>, 104(A4), 6729-6736. doi: https://doi.org/10 .1029/1998JA900183 Yang, Y., Shen, F., Yang, Z., & Feng, X. (2018). Prediction of solar wind speed at 1 au using an artificial neural network. <i>Space Weather</i>, 16(9), 1227-1244. doi: