

Probabilistic Geomagnetic Storm Forecasting via Deep Learning

Adrian Tasistro-Hart^{1,2}, Alexander Grayver¹, Alexey Kuvshinov¹

¹Institute of Geophysics, ETH Zürich, Sonneggstrasse 5, 8092 Zürich, Switzerland

²Department of Earth Science, University of California, Santa Barbara, CA 93106, USA

Key Points:

- We present a neural network architecture that utilizes both observations from the L1 point and solar disk, improving forecast reliability
- Our neural network architecture learns reliable estimates of uncertainty in multiple hour ahead forecasts
- Instead of the conventional disturbance storm time (Dst) index we forecast the external component of geomagnetic storms, Est

Corresponding author: Adrian Tasistro-Hart, adrian.tasistro-hart@ucsb.edu

Abstract

Geomagnetic storms, which are governed by the plasma magnetohydrodynamics of the solar-interplanetary-magnetosphere system, entail a formidable challenge for physical forward modeling. Yet, the abundance of high quality observational data has been amenable for the application of data-hungry neural networks to geomagnetic storm forecasting. Previous applications of neural networks to storm forecasting have utilized solar wind observations from the Earth-Sun first Lagrangian point (L1) or closer and have all generated deterministic output without uncertainty estimates. Furthermore, forecasting work has focused on indices that are also sensitive to induced internal magnetic fields, complicating the forecasting problem with another layer of non-linearity. We address these points, presenting neural networks trained on observations from both the solar disk and the L1 point. Our architecture generates reliable probabilistic forecasts over Est, the external component of the disturbance storm time index, showing that neural networks can gauge confidence in their output.

Plain Language Summary

Geomagnetic storms are capable of damaging infrastructure like power grids and communication lines, motivating our need to forecast them. Solar phenomena produce geomagnetic storms, which occur when these phenomena reach Earth as bursts of the solar wind. Decades of satellite observations of both the solar wind near the Earth and of the Sun itself are promising for forecasting geomagnetic storms with algorithms known as neural networks. Several neural network architectures have been applied to geomagnetic storm forecasting, but their full potential remains unexplored. First, all existing neural networks have used measurements of the solar wind one hour upstream of the Earth or closer. While these observations are critical for understanding geomagnetic storm progression, from them it is nearly impossible to forecast more than an hour in advance. We include observations of the Sun itself, which reach Earth much faster than the solar wind, thereby including information for forecasting further in advance. Second, all existing neural networks have generated forecasts without uncertainty estimates, meaning that end-users (such as utilities or telecommunications companies) know little about forecast confidence. We present an architecture that generates estimates of uncertainty, and our results demonstrate that neural networks learn how confident to be in their forecasts.

44 **1 Introduction**

45 Mankind has experienced a number of blackouts caused by geomagnetically induced
46 currents (GICs), which can result in millions of dollars of damages and leave millions with-
47 out electricity (Bolduc, 2002; Love et al., 2018). The possibility of such disruptions has
48 motivated the goal of forecasting GICs. All GICs in turn result from geomagnetic storms,
49 which generate the variation in Earth’s external field that induces GICs. The problem
50 of forecasting GICs then amounts to forecasting geomagnetic storms. These storms re-
51 sult from the propagation of solar activity via the solar wind and its coupling to Earth’s
52 magnetosphere. Given abundant observational data of the solar wind and disk as well
53 as of Earth’s magnetic field, the application of data-hungry deep learning algorithms is
54 suitable for the forecasting problem.

55 **1.1 Geomagnetic storms**

56 Geomagnetic storms have traditionally been quantified by indices such as the dis-
57 turbance storm time (Dst in nT) or Kp (unitless) indices (e.g. Bartels et al. (1939)), both
58 of which register deviations from the quiet time horizontal component of Earth’s mag-
59 netic field. The basic mechanism of geomagnetic storm formation is the strengthening
60 of Earth’s ring current in response to changing solar wind conditions, and this strength-
61 ened current system generates a magnetic field that counters Earth’s dipole, weakening
62 it relative to quiet conditions (Daglis et al., 1999). The solar wind parameters most im-
63 portant for strengthening the ring current are its southward component of the inter-planetary
64 magnetic field (IMF), velocity, and plasma density, which all positively impact storm am-
65 plitude (Daglis et al., 1999; Gonzalez et al., 1999; Wolf et al., 1997). All solar wind ac-
66 tivity that generates significant, rapid fluctuation in Earth’s external magnetic field poses
67 a threat to ground-based conducting systems, such as power and communication lines,
68 during geomagnetic storms.

69 **1.2 Why deep learning?**

70 Given the complexity of the underlying physics, which involves the magnetohydro-
71 dynamics (MHD) and plasma physics of propagating solar activity through the solar wind
72 and its subsequent interaction with Earth’s magnetosphere, a fully physical forward model
73 of the system would be both computationally expensive and poorly constrained. At the

74 same time, given that we are aware of the important physical quantities responsible for
75 geomagnetic storms, such a physical model is overkill for the problem of forecasting the
76 low-order response of Earth’s magnetic field to solar activity.

77 For this reason, the first approach to geomagnetic storm modeling took the form
78 of simple empirical models that related the time rate of change of Dst to solar wind pa-
79 rameters. The pioneering work was a three-term deterministic model developed by Bur-
80 ton et al. (1975), but its simplicity, while elegant, often generates inaccurate forecasts.
81 Subsequent modeling has attempted to improve accuracy by adding more degrees of free-
82 dom. For example, while obtaining more predictive power, Temerin & Li (2006) added
83 almost a dozen more terms with significantly more complex functional forms, sacrific-
84 ing the simplicity of the initial model. Neural networks (NNs), which form the backbone
85 of deep learning, are the logical conclusion to the exercise of adding more and more heuris-
86 tic functional forms, since the task of a NN is to learn the relevant functions rather than
87 have them prescribed (Leshno et al., 1993; Pinkus, 1999). While NNs provide a flexible
88 framework for learning the appropriate functions that dictate the system dynamics, it
89 is more difficult to interpret the complex interactions of neurons and the abundance of
90 learnable parameters (weights and biases). The latter also means that NNs require abun-
91 dant data to be effectively trained. Fortunately, decades of observations of the solar wind
92 and solar disk provide a suitable dataset necessary for deep learning.

93 **1.3 Prior applications of deep learning to geomagnetic storm forecast-** 94 **ing**

95 Previous work with NNs has focused almost entirely on prediction of Dst or other
96 indices of geomagnetic activity, such as the Kp and the auroral electrojet (AE) indices.
97 Supplemental Table S1 provides a succinct review of the application of NNs to the fore-
98 casting of Dst (Andriyas & Andriyas, 2015; Bala & Reiff, 2012; Gleisner et al., 1996; Jankovičová
99 et al., 2002; Kugblenu et al., 1999; Lazzús et al., 2017; Munsami, 2000; Pallocchia et al.,
100 2006; Revallo et al., 2014; Sharifie et al., 2006; Stepanova et al., 2005; Stepanova & Pérez,
101 2000; Wei et al., 2007; Wu & Lundstedt, 1996, 1997). These studies have applied a va-
102 riety of architectures and data sources, but in generating forecasts for Dst, most have
103 used the basic solar wind parameter measurements as well as prior values of Dst. All pre-
104 vious studies applying NNs to Dst forecasting to our knowledge have utilized observa-
105 tions made at the Earth-Sun L1 point or closer. Furthermore, all studies to date using

106 NNs to forecast Dst (or any other geomagnetic storm index) have been deterministic,
 107 generating predictions without any measure of uncertainty. Only Gruet et al. (2018) at-
 108 tempt to assess uncertainty in their forecasts, but they use a Gaussian process model to
 109 develop uncertainty estimates, and their NN output is deterministic.

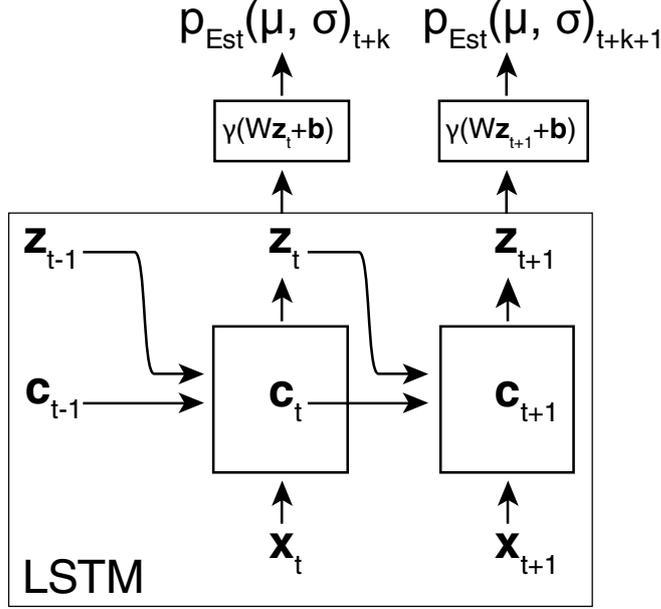
110 This work improves on previous advances by presenting the first application of prob-
 111 abilistic neural networks to geomagnetic storm forecasting, utilizing recent developments
 112 in deep learning to generate forecasts with uncertainty. Our networks are capable of learn-
 113 ing how confident to be in their predictions, and in doing so improve forecast reliabil-
 114 ity. These networks take as input not just observations from the L1 point but also ob-
 115 servations of radiative phenomena on the solar disk. Finally, instead of forecasting Dst,
 116 we focus on its external component, Est, which does not incorporate the effects of Earth’s
 117 subsurface conductivity structure.

118 2 Data and Methods

119 2.1 Probabilistic Neural Network Architecture

120 Recently, recurrent architectures for time series regression have emerged that com-
 121 bine ridge functions with state vectors to create units with “memory”. The most suc-
 122 cessful of these has been the long short-term memory architecture (LSTM), introduced
 123 by Hochreiter & Schmidhuber (1997). The LSTM cell, as its name implies, uses new in-
 124 put data with both the previous output and previous internal state to update its inter-
 125 nal state and generate new output (Supplement, Text S6). This architecture has been
 126 applied to Dst forecasting by Gruet et al. (2018), but, like all previous applications of
 127 neural networks to storm forecasting (summarized in the Supplement, Table S1), the net-
 128 work generated deterministic output with no prediction of forecast uncertainty.

133 We present an architecture (Figure 1) in which the NN learns to assess uncertainty
 134 in its own forecast, thereby generating probabilistic forecasts. The two basic layers uti-
 135 lized within this architecture are LSTM and dense layers. The former is described above,
 136 and the latter is an implementation of the so-called “fully connected hidden layer”, which
 137 references the fact that each entry in \mathbf{z} depends on all of the outputs from the preced-
 138 ing layer via W . That is, a dense layer that receives inputs $\mathbf{x} \in \mathbb{R}^n$ from a preceding
 139 network layer in turn generates an output vector $\mathbf{z} \in \mathbb{R}^m$ via the operation $\mathbf{z} = \gamma(W\mathbf{x} +$
 140 $\mathbf{b})$ with $W \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$, where n is the dimensionality of the preceding hid-



129 **Figure 1.** Schematic architecture for the deterministic network that learns parameters over
 130 an output distribution for Est. Our output distribution is a Gaussian. Two full time steps of
 131 network iteration are shown, with the portion of the network enclosing the LSTM cell labeled
 132 “LSTM”.

141 den layer, m is the dimensionality of the current hidden layer, and $\gamma(\cdot)$ is a nonlinear ac-
 142 tivation function that acts element-wise.

143 Inputs into our NN architecture are fed directly to an LSTM cell, and outputs from
 144 the LSTM cell are fed through a series of fully-connected hidden layers. The outputs from
 145 the last hidden layer are parameters for an output distribution over Est. We choose to
 146 use a Gaussian output distribution and compare other alternatives in the Supplement
 147 (Text S3).

The simplest cost function in this probabilistic framework is precisely the output
 distribution itself evaluated as a likelihood of observed data y (i.e., Est at some time $t+$
 k) with respect to the distribution parameters generated from the given input:

$$C(\mathbf{x}, y) = -\log p(y|\mu(W, \mathbf{b}, \mathbf{x}), \sigma(W, \mathbf{b}, \mathbf{x})) \quad (1)$$

148 where the distribution parameters μ and σ depend on the network weights and biases
 149 and can thus be differentiated against them. However, when learning two-parameter dis-
 150 tributions, the parameter for scale often introduces leniency in the output distribution,

151 allowing the network to expand uncertainty in its forecast rather than move its estimate
 152 for the center (see supplement Text S3).

153 We found that utilizing a Gaussian output distribution with a regularized Gaus-
 154 sian likelihood as the cost function performed well for geomagnetic storm forecasting.
 155 Equation 2 shows the form for this modified log-likelihood in which $\alpha (y - \mu)^2 + \beta \frac{1}{\sigma^2}$
 156 are the terms that we have added, introducing α and β as additional hyper-parameters.
 157 This regularization encourages the network to learn more reasonable estimates for μ , off-
 158 setting the normalization by σ^2 , while also allowing the user to further incentivize ($\beta >$
 159 0) or penalize ($\beta < 0$) expanding forecast uncertainty.

$$C_{\text{Gaussian, regularized}}(y, \mu, \sigma) = \log(\sqrt{2\pi}\sigma) + \frac{(y - \mu)^2}{2\sigma^2} + \alpha (y - \mu)^2 + \beta \frac{1}{\sigma^2} \quad (2)$$

160 Other approaches have been formulated for learning uncertainty via neural networks, such
 161 as Bayes-by-Backprop (Blundell et al., 2015), which represents uncertainty in the net-
 162 work weights rather than in its output. Our implementation of this approach was not
 163 useful for storm forecasting (Supplement Text S2).

164 For all implementations, training, and testing of neural networks, we use Python
 165 wrappers for the learning framework TensorFlow (Abadi et al., 2015). This framework
 166 is capable of representing neurons and the functional operations relating them as well
 167 as numerically computing the relevant gradients to train the network. TensorFlow pro-
 168 vides an implementation of the high level deep learning Keras API (<https://keras.io/>),
 169 which allows for modular construction of networks from the layers described above. We
 170 also make use of the recently introduced TensorFlow Probability library, which provides
 171 a straightforward means of adding probability distributions as layers, allowing outputs
 172 from previous layers to be used as parameters for the distribution layers. These layers
 173 are compiled into a model that contains all the operations of the entire network as well
 174 as the particular cost functions and optimizers that dictate the learning process for given
 175 training inputs and outputs. We use the Adam optimizer for gradient update steps (Kingma
 176 & Ba, 2014). All neural network configuration and training parameters are listed in Sup-
 177 plement Text S4.

178

2.2 Output Data

179

180

181

182

Most geomagnetic storm forecasting thus far has emphasized prediction of Dst. However, Dst is actually a sum of internal and external components (Equation 3), and the internal component, Ist, is generated by currents induced in Earth's subsurface by variation in the external component, Est (Maus & Weidelt, 2004) (Equation 4),

$$Dst(t) = Ist(t) + Est(t) \quad (3)$$

$$Ist(t) = \int_{-\infty}^t Q(t - \tau) Est(\tau) d\tau, \quad (4)$$

183

184

where $Q(t-\tau)$ is the induction kernel that depends on a radial electrical conductivity profile of subsurface.

185

186

187

188

189

190

191

192

The problem of forecasting Dst is then actually two separate problems: the first is forecasting Est, and the second is learning Earth's induction response, Ist. With a suitable model of Earth's subsurface conductivity structure, however, knowledge of Est is sufficient to reconstruct Ist and thereby Dst. Furthermore, because the external field is what responds to magnetospheric activity anyway, it is more natural to forecast Est than Dst. Therefore, we generate forecasts of Est, and the data accessed from NOAA were generated following the methodology of Maus & Weidelt (2004), who utilize a one-dimensional conductivity model for the decomposition of Dst into Ist and Est.

193

194

195

196

197

198

199

200

201

202

This approach will be increasingly important as we attempt to forecast higher order structure in Earth's external field. Est captures only the first zonal (dipole) component of external magnetic field variability, but significant variation exists on shorter spatial scales, where the interaction with local conductivity structures becomes more important and complicated (Kelbert, 2020). Given that the ultimate goal of geomagnetic storm forecasting is to forecast GICs, it is important to note that GICs themselves depend strongly on local conductivity structures and local external magnetic field variability (Olsen & Kuvshinov, 2004; Pütke et al., 2014). The first step to forecasting these higher order external field coefficients is forecasting a single external field coefficient, Est, which is the focus of this work.

203

2.3 Input Data

204

205

Two basic observation types relevant to geomagnetic storm forecasting have been made for the past few decades: the first includes measurements of the solar wind made

206 in situ around the L1 point in the Earth-Sun system, and the second are measurements
207 made directly of the solar disk and corona. These two data streams provide related but
208 temporally disjoint information. Radiative phenomena on the solar disk take under nine
209 minutes to be observed at Earth, while the solar wind requires two to five days to prop-
210 agate from the solar disk to the L1 point. The L1 point is only 1.5×10^6 km from Earth,
211 however, which is approximately one hour travel time at typical solar wind speeds (mean
212 solar wind speed is roughly 440 km s^{-1} from the OMNI dataset).

213 Thus, while solar wind measurements near Earth's magnetosphere are ultimately
214 the most relevant quantities for accurate geomagnetic storm forecasting, using only ob-
215 servations from the L1 point limits the forecast time to roughly an hour (Shprits et al.,
216 2019). On the other hand, solar activity is ultimately responsible for all geomagnetic storms,
217 but identifying which phenomena on the solar disk have the potential to cause geomag-
218 netic storms and predicting the storm lag times and amplitudes resulting from those phe-
219 nomena are not trivial tasks. Observations from the solar disk include measurements of
220 coronal mass ejections (CMEs) around the perimeter of the disk, images of the solar sur-
221 face at various wavelengths, and surface radiative fluxes.

222 Input data come from three sources: the OMNI, GOES, and SOHO LASCO CME
223 datasets. All details on data preprocessing are briefly discussed in the supplement (Text
224 S2).

225 The low (hourly) resolution OMNI data include several solar wind and solar ob-
226 servations, which we extracted for the years 1995-2018. During this time interval, mea-
227 surements of the solar wind (SW) come from the ACE and DSCVR satellite missions
228 at the L1 point. The quantities that we use as input data from this dataset are the three
229 components of the interplanetary magnetic field in geocentric solar magnetospheric (GSM)
230 coordinates, SW velocity, SW particle density, SW temperature, and SW longitude and
231 latitude incident on Earth's magnetosphere.

232 The GOES mission provided two time series of x-ray fluxes integrated over the so-
233 lar disk. One series covered the wavelengths from $0.5\text{-}4 \text{ \AA}$, and the other series covered
234 wavelengths $1\text{-}8 \text{ \AA}$. Given that these measurements vary over orders of magnitude, we
235 take their logarithm as input. These data were reliably available from 1986 onwards.

236 The LASCO SOHO CME database provides a catalogue of CMEs observed around
237 the perimeter of the solar disk, with five basic quantities estimated for each event: cen-
238 tral position angle, angular width, speed, mass, and kinetic energy. Three estimates of
239 speed are reported in the catalogue, all of which we include as training input. We only
240 consider CMEs for which all data fields are reported, and during hours with multiple events,
241 we take only the event with the largest estimated kinetic energy. The estimates of mass
242 and kinetic energy varied over orders of magnitude, so we instead take their logarithms
243 as input. The database contains measurements from 1996-2018.

244 Finally, previous observations of Est were included as input while forecasting fu-
245 ture values. In total then data is available roughly from 1996 to 2018. Of these 22 years,
246 we take 18 years (approximately 158,000 hours) as training data, and 4 years for test-
247 ing data (approximately 35,000 hours).

248 **2.4 Evaluating network performance**

249 In most geomagnetic storm forecasting to date, forecast accuracy is assessed by met-
250 rics such as the root mean square error and the Pearson correlation coefficient between
251 forecasted and observed data. However, these statistics are generally not compared to
252 those of a null hypothesis, for example persistence forecasting in which the best estimate
253 for any time in the future is simply the last observed value. Due to the auto-correlative
254 nature of the Est time series, persistence forecasting actually generates deceptively high
255 correlations and low errors (Supplement Figure S8) (Shprits et al., 2019). In fact, all the
256 networks in Supplement Table S1 either underperform or barely outperform persistence
257 forecasting as quantified by these two metrics. Furthermore, these metrics are computed
258 over both quiet and disturbed times, while the ability to predict storms is the task of
259 interest. Other than refining consideration of these metrics to only storm main phases,
260 it is not obvious what metric best evaluates forecasting performance for models with de-
261 terministic output.

262 With probabilistic networks, however, reliability curves provide a useful and eas-
263 ily interpretable metric to evaluate forecast performance. Each curve corresponds to a
264 threshold in Est, and the axes compare the observed probability of exceeding that thresh-
265 old compared to the predicted probability. A perfectly reliable forecast would generate
266 curves that fall on the 1:1 line through the origin for all thresholds. If the observed re-

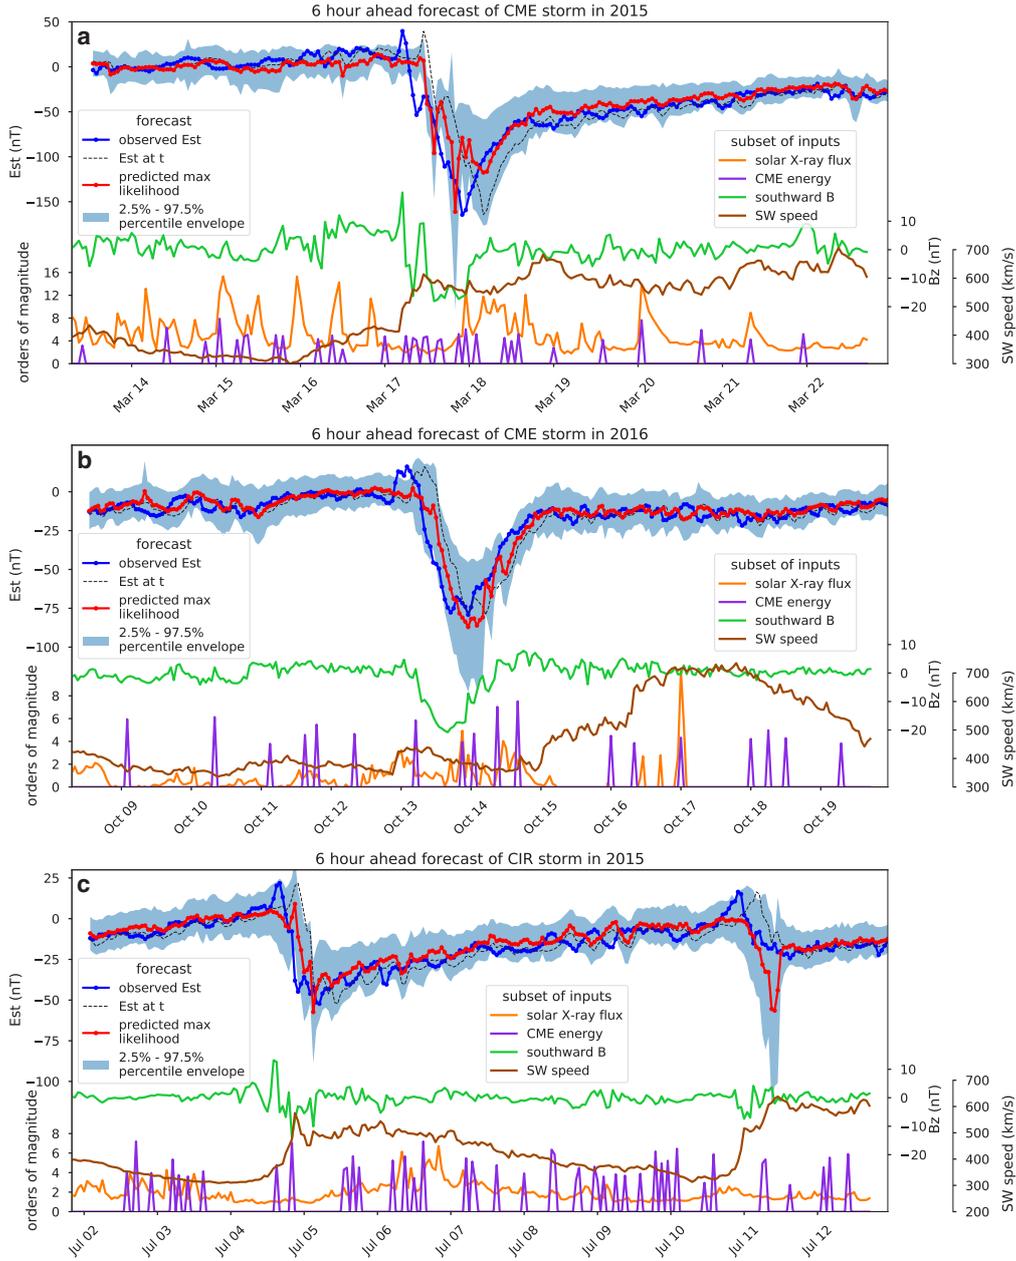
267 liability curve plots over the 1:1 line, then the forecast underestimates the occurrence
268 of events exceeding that threshold, while if the reliability curve plots under the 1:1 line,
269 the forecast is conservative and overestimates the occurrence of events exceeding the given
270 threshold. Because these statistics are computed for thresholds in Est, the reliability as-
271 sessment method by construction evaluates storm time forecasting separately from quiet
272 time forecasting. We utilize reliability curves to assess forecast performance.

273 Computing reliability curves requires binning data by intervals of predicted exceedance
274 probabilities, which means that empirical statistics for infrequent, large storms will be
275 less well constrained than smaller storms, particularly at large forecast probabilities. To
276 assess uncertainty in the reliability curve computation, we use bootstrap resampling of
277 forecasted and observed threshold exceedances to compute confidence intervals over ob-
278 served threshold exceedances for a given bin of predicted threshold exceedance. Further-
279 more, we compute consistency intervals that indicate for the amount of data in each bin
280 the spread in forecasting skill that one might anticipate from a perfectly reliable fore-
281 cast (Bröcker & Smith, 2007).

282 **3 Results and Discussion**

283 In general, our architecture is capable of learning meaningful and reliable measures
284 of uncertainty in its forecasts. We discuss the architecture performance here by restrict-
285 ing ourselves to six-hour ahead probabilistic forecasts from the network trained with both
286 L1 and solar disk inputs for three different storms: two caused by CMEs (Figure 2A &
287 B) and one resulting from a co-rotating interaction region (CIR) (Figure 2C). One to
288 six hour ahead forecasts, as well as detailed comparison of the networks trained with only
289 L1 data and both L1 and solar disk inputs, can be found in the Supplement.

298 The first storm, in March of 2015, provides a prototypical example of a geomag-
299 netic storm caused by a magnetic cloud emitted by the mass ejection visible as the spike
300 in CME energy at the beginning of March 15 (Figure 2A, note that the axis is orders
301 of magnitude) (Patel et al., 2019). The energetic mass ejection is associated with a peak
302 in integrated x-ray fluxes, followed two days later by a relatively large geomagnetic storm
303 beginning on March 17. The storm main phase is associated with a sustained southward
304 IMF of roughly -20 nT and roughly doubled solar wind speeds. In this situation, given
305 the clear connection between the mass ejection and the ensuing storm, we would expect



290 **Figure 2.** Six-hour-ahead probabilistic forecasts on testing data for two CME storms and
 291 one CIR storm, as identified by Patel et al. (2019); Shen et al. (2017). Order of magnitude vari-
 292 ability in x-ray flux (long channel is plotted) is exaggerated by ten times. The black dashed line
 293 for Est is a persistence forecast. **(a)** Severe CME geomagnetic storm (min Dst=-222 nT). Large
 294 amplitude variability in the southward component of the IMF might be responsible for fore-
 295 cast variability as the storm enters the main phase. **(b)** Intense CME geomagnetic storm (min
 296 Dst=-104 nT). **(c)** Intense CIR geomagnetic storm (min SYM-H=-87 nT). Order of magnitude
 297 variability in x-ray flux here is only exaggerated by five times.

306 a successfully trained network to be able to expand uncertainty in its forecast as con-
307 ceivable storm arrival times approach, reflecting an understanding of the causal associ-
308 ation between activity on the solar disk and geomagnetic storms. Yet, forecast uncer-
309 tainties only expand when disturbed solar wind reaches the L1 point. At that time, the
310 network becomes aware of the storm arrival and adjusts its output by dropping Est fore-
311 casts and increasing forecast uncertainty (Figure 2A). The same is true for the smaller
312 amplitude CME storm of October 13, 2016 in Figure 2B, where forecast uncertainty only
313 grows as soon as the storm arrives at the L1 point. This storm is associated with the CME
314 visible on October 10 (Patel et al., 2019), so the occurrence of other CMEs of similar mag-
315 nitude (e.g. on October 9 and 11) demonstrates non-uniqueness that illustrates why the
316 network struggles to identify geoeffective solar activity from the provided inputs.

317 The final storm on July 4, 2015 was chosen because it corresponds to a CIR (Shen
318 et al., 2017), as evidenced by a lack of sustained, southward IMF, a step increase in so-
319 lar wind velocity, and relatively low amplitude storm-time Est (Figure 2C). The nature
320 of CIR storms differs from those originating from CMEs (Zhang et al., 2007), so we sought
321 to investigate if the forecast for CIR storms differs from that for CME storms. Again,
322 for the storm on July 4, the network is unable to preemptively expand forecast uncer-
323 tainties in response to information from solar disk, demonstrating that inputs from the
324 L1 point dominate the forecast. On July 11, the network mistakenly forecasts a storm
325 main phase, likely in response to the increased solar wind speed that did not actually
326 generate a substantial main phase.

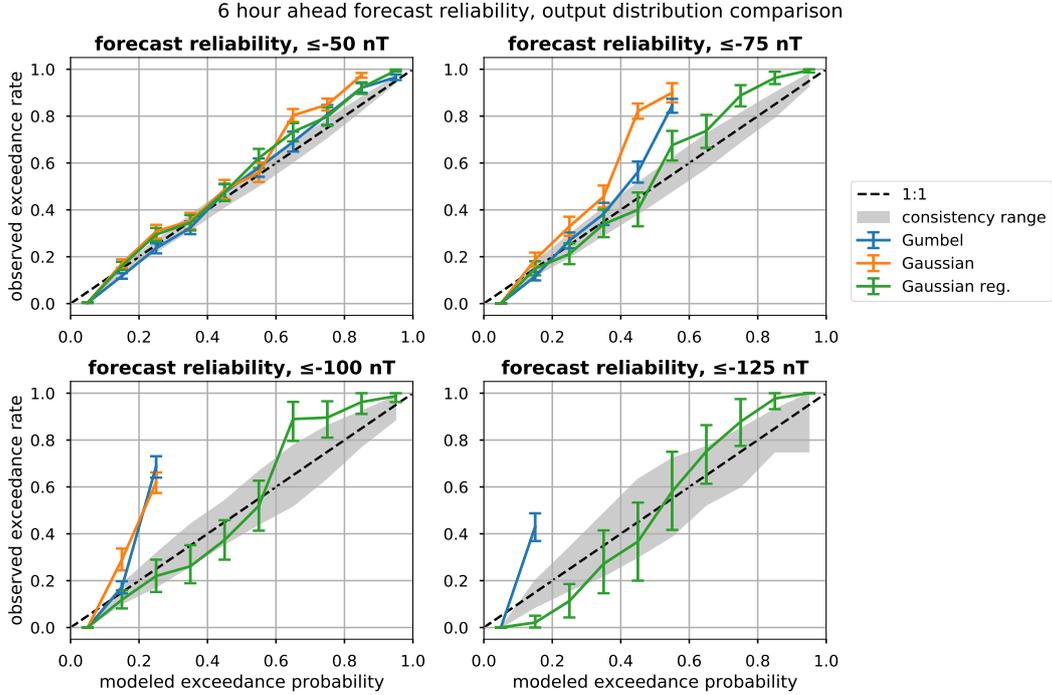
327 In all cases, forecasting geomagnetic storm onset is not improved by utilizing ob-
328 servational inputs from the solar disk, but recovery is generally well-predicted, and fore-
329 casts deviate from persistence, meaning that the network is not just taking the last ob-
330 served Est value for its next forecast. Unlike previous results, our network is capable of
331 generating meaningful estimates of uncertainty in its forecasts. In all cases, once the net-
332 work detects the possible onset of a geomagnetic storm, it expands its forecast uncer-
333 tainty, generally maintaining observed Est values within the 95% forecast confidence in-
334 terval and providing reliable multiple hour ahead forecasts (see Supplement Text S5 for
335 one to six hour ahead forecasts). After storm main phases, forecast uncertainties decrease
336 during the generally well-predicted recovery phase. Given that the recovery phase is dic-
337 tated by the internal dynamics of the ring current decay (and thereby independent of
338 the solar wind state) (Daglis, 2007), its predictability is reasonable. Thus, our network

339 exhibits forecast uncertainties that are consistent with where one would anticipate the
340 greatest uncertainty in geomagnetic storm development with information from the L1
341 point, namely, the storm onset and main phase.

342 Reliabilities for four different storm thresholds generally overlap with the consis-
343 tency intervals for each bin, demonstrating that our network generates reliable forecasts
344 (Figure 3). For threshold of -75 and -100 nT, forecasted exceedance probabilities in the
345 range of 0.7-0.9 tended to slightly underestimate observed exceedance rates, which is con-
346 sistent with the observation that storm onsets remain difficult to predict exactly.

347 Notably, the regularization of the cost function for a Gaussian output distribution
348 significantly improves forecast reliability. Networks trained with unregularized Gaussian
349 and Gumbel output distributions (Supplement Text S3) are unable move the location
350 parameters of their forecasts during large amplitude storms, preferring instead to expand
351 forecast uncertainty, meaning that peak storm times, while often within the 95% con-
352 fidence interval, are only predicted at extremely low exceedance probabilities. This be-
353 havior explains why the reliability curves lack data to bin at high forecast probabilities
354 and furthermore why storms are underestimated for lower exceedance probabilities (Fig-
355 ure 3).

365 While some improvement in forecast reliability for smaller magnitude storms (Est
366 thresholds of -50 and -75 nT) does seem to result from the incorporation of data from
367 the solar disk (Supplement Figure S7), the preceding discussion and the result that fore-
368 cast behavior does not qualitatively change by adding solar disk inputs (Supplement Text
369 S5) indicates that we are unable to successfully utilize observations from the solar disk
370 to forecast storm arrival and amplitude. This shortcoming suggests that the informa-
371 tion necessary for identifying geoeffective solar activity is lacking in the training data,
372 and/or that the network architecture is inadequate for utilizing these data. For instance,
373 the x-ray fluxes are integrated over the entire solar disk, but peaks in these fluxes can
374 often be associated with flare events, which themselves often occur simultaneously with
375 geoeffective mass ejections (Tobiska et al., 2013). Larger, more central flares are asso-
376 ciated with larger geomagnetic storms, so adding time series of flare occurrences with
377 locations on the solar disk would complement the input series of x-ray fluxes and CMEs
378 (Tobiska et al., 2013). Furthermore, the CME dataset only includes ejections visible around
379 the rim of the solar disk, while geoeffective ejections occur towards the center. Thus, only



356 **Figure 3.** Reliability curves for the networks with a Gumbel output and cost (blue), Gaus-
 357 sian output and cost (orange), and Gaussian output with regularized cost (green). All curves
 358 are for a six hour ahead forecast for four Est thresholds in eleven bins. Exceedance was taken in
 359 the negative sense, i.e., taking values less than or equal to the given threshold. Error bars show
 360 the 2.5-97.5 percentile range from bootstrapped resampling (number of bootstrapped samples
 361 was 1000) within the bins of forecasted exceedance probability. The envelopes show the 2.5-97.5
 362 percentile range from consistency resampling of a perfectly reliable forecast, demonstrating the
 363 conceivable range in reliable forecasts given the number of data in each bin (Bröcker & Smith,
 364 2007). The regularized Gaussian network is the most reliable of the three.

380 centralized ejections that also emit an observable lobe beyond the rim of the solar disk
381 could be reliably associated with geomagnetic storms, potentially rendering the CME
382 database largely irrelevant for the problem of geomagnetic storm forecasting. Finally,
383 integrated solar x-ray flux peaks from flares have been empirically related to solar wind
384 speeds and geomagnetic storm amplitudes, thereby providing a means of learning lag times
385 between solar activity and storm arrivals (Tobiska et al., 2013). However, LSTM net-
386 works struggle with learning lag times (Gers et al., 2002), so the network architecture
387 we have utilized is not amenable to this task.

388 4 Conclusions

389 This work has demonstrated a NN architecture capable of learning reliable mea-
390 sures of uncertainty in its forecasts of geomagnetic storms. Learning uncertainty in NN
391 output results in more useful probabilistic forecasts than learning uncertainty in the NN
392 parameters, and the choice of output distribution and cost function has a large impact
393 on the resulting reliability of the trained network. Specifically, adding regularizing terms
394 in the likelihood cost function improves the forecast reliability by incentivizing networks
395 to forecast more reasonable mean values rather than simply increasing forecast uncer-
396 tainty.

397 These neural networks are also the first to utilize as inputs observations from the
398 solar disk and L1 point. This provides improved uncertainty forecast and higher reli-
399 ability, although storm arrival and amplitude forecasting did not substantially improve
400 from the inclusion of these data. Thus, leveraging time series of observations of the so-
401 lar disk, which are often sparse, remains an open problem, and future network architec-
402 tures must be carefully designed to utilize these data sources.

403 Acknowledgments

404 This work was partially supported by the ESA through the Swarm DISC project. The
405 low (hourly) resolution OMNI data were obtained from the GSFC/SPDF OMNIWeb in-
406 terface (<https://omniweb.gsfc.nasa.gov>). We acknowledge use of the LASCO SOHO
407 CME catalog (https://cdaw.gsfc.nasa.gov/CME_list/), which is generated and main-
408 tained at the CDAW Data Center by NASA and The Catholic University of America in
409 cooperation with the Naval Research Laboratory. SOHO is a project of international co-
410 operation between ESA and NASA. We acknowledge use of GOES mission x-ray flux data,

411 accessed from <https://satdat.ngdc.noaa.gov/sem/goes/data/avg/>. Finally, we ac-
412 knowledge use of the EST-IST-DST dataset from NOAA accessed from <https://www>
413 [.ngdc.noaa.gov/geomag/est_ist.shtml](https://www.ngdc.noaa.gov/geomag/est_ist.shtml). All data and analysis presented in this study
414 are available as Jupyter notebooks at <https://doi.org/10.5281/zenodo.3751682>.

415 References

- 416 Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng,
417 X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*.
418 Retrieved from <https://www.tensorflow.org/> (Software available from tensor-
419 flow.org)
- 420 Andriyas, T., & Andriyas, S. (2015). Relevance vector machines as a tool for fore-
421 casting geomagnetic storms during years 1996–2007. *Journal of Atmospheric and*
422 *Solar-Terrestrial Physics*, 125, 10–20.
- 423 Bala, R., & Reiff, P. (2012). Improvements in short-term forecasting of geomagnetic
424 activity. *Space Weather*, 10(6).
- 425 Bartels, J., Heck, N. H., & Johnston, H. F. (1939). The three-hour-range index mea-
426 suring geomagnetic activity. *Journal of Geophysical Research*, 44(4), 411. Re-
427 trieved from <http://doi.wiley.com/10.1029/TE044i004p00411> doi: 10.1029/
428 TE044i004p00411
- 429 Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015). Weight uncer-
430 tainty in neural networks. *arXiv preprint arXiv:1505.05424*.
- 431 Bolduc, L. (2002). GIC observations and studies in the hydro-Québec power
432 system. *Journal of Atmospheric and Solar-Terrestrial Physics*, 64(16), 1793 -
433 1802. Retrieved from [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S1364682602001281)
434 [S1364682602001281](http://www.sciencedirect.com/science/article/pii/S1364682602001281) (Space Weather Effects on Technological Systems) doi:
435 [https://doi.org/10.1016/S1364-6826\(02\)00128-1](https://doi.org/10.1016/S1364-6826(02)00128-1)
- 436 Bröcker, J., & Smith, L. A. (2007, jun). Increasing the Reliability of Reliability
437 Diagrams. *Weather and Forecasting*, 22(3), 651–661. Retrieved from [http://](http://journals.ametsoc.org/doi/abs/10.1175/WAF993.1)
438 journals.ametsoc.org/doi/abs/10.1175/WAF993.1 doi: 10.1175/WAF993.1
- 439 Burton, R. K., McPherron, R., & Russell, C. (1975). An empirical relationship be-
440 tween interplanetary conditions and Dst. *Journal of geophysical research*, 80(31),
441 4204–4214.
- 442 Daglis, I. A. (2007, feb). Ring Current Dynamics. *Space Science Reviews*, 124(1-4),

- 443 183–202. Retrieved from <http://link.springer.com/10.1007/s11214-006-9104>
444 -z doi: 10.1007/s11214-006-9104-z
- 445 Daglis, I. A., Thorne, R. M., Baumjohann, W., & Orsini, S. (1999). The terres-
446 trial ring current: Origin, formation, and decay. *Reviews of Geophysics*, *37*(4),
447 407–438.
- 448 Gers, F. A., Schraudolph, N. N., & Schmidhuber, J. (2002). Learning precise tim-
449 ing with LSTM recurrent networks. *Journal of Machine Learning Research*, *3*(1),
450 115–143. doi: 10.1162/153244303768966139
- 451 Gleisner, H., Lundstedt, H., & Wintoft, P. (1996). Predicting geomagnetic storms
452 from solar-wind data using time-delay neural networks. *Annales Geophysicae*, *14*,
453 679.
- 454 Gonzalez, W. D., Tsurutani, B. T., & De Gonzalez, A. L. C. (1999). Interplanetary
455 origin of geomagnetic storms. *Space Science Reviews*, *88*(3-4), 529-562.
- 456 Gruet, M., Chandorkar, M., Sicard, A., & Camporeale, E. (2018). Multiple-hour-
457 ahead forecast of the Dst index using a combination of long short-term memory
458 neural network and Gaussian process. *Space Weather*, *16*(11), 1882–1896.
- 459 Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural compu-*
460 *tation*, *9*(8), 1735–1780.
- 461 Jankovičová, D., Dolinský, P., Valach, F., & Vörös, Z. (2002). Neural network-
462 based nonlinear prediction of magnetic storms. *Journal of atmospheric and solar-*
463 *terrestrial physics*, *64*(5-6), 651–656.
- 464 Kelbert, A. (2020). *The Role of Global/Regional Earth Conductivity Models in*
465 *Natural Geomagnetic Hazard Mitigation* (Vol. 41) (No. 1). Springer Netherlands.
466 Retrieved from <https://doi.org/10.1007/s10712-019-09579-z> doi: 10.1007/
467 s10712-019-09579-z
- 468 Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*
469 *preprint arXiv:1412.6980*.
- 470 Kugblenu, S., Taguchi, S., & Okuzawa, T. (1999). Prediction of the geomagnetic
471 storm associated Dst index using an artificial neural network algorithm. *Earth*
472 *Planets Space*, *51*(307-313).
- 473 Lazzús, J., Vega, P., Rojas, P., & Salfate, I. (2017). Forecasting the Dst index using
474 a swarm-optimized neural network. *Space Weather*, *15*(8), 1068–1089.
- 475 Leshno, M., Lin, V. Y., Pinkus, A., & Schocken, S. (1993). Multilayer feed-

- 476 forward networks with a nonpolynomial activation function can approximate
477 any function. *Neural Networks*, 6(6), 861 - 867. Retrieved from [http://](http://www.sciencedirect.com/science/article/pii/S0893608005801315)
478 www.sciencedirect.com/science/article/pii/S0893608005801315 doi:
479 [https://doi.org/10.1016/S0893-6080\(05\)80131-5](https://doi.org/10.1016/S0893-6080(05)80131-5)
- 480 Love, J. J., Lucas, G. M., Kelbert, A., & Bedrosian, P. A. (2018). Geoelectric
481 Hazard Maps for the Mid-Atlantic United States: 100 Year Extreme Values
482 and the 1989 Magnetic Storm. *Geophysical Research Letters*, 45(1), 5–14. doi:
483 10.1002/2017GL076042
- 484 Maus, S., & Weidelt, P. (2004). Separating the magnetospheric disturbance mag-
485 netic field into external and transient internal contributions using a 1D con-
486 ductivity model of the Earth. *Geophysical Research Letters*, 31(12), 2–5. doi:
487 10.1029/2004GL020232
- 488 Munsami, V. (2000). Determination of the effects of substorms on the storm-
489 time ring current using neural networks. *Journal of Geophysical Research: Space*
490 *Physics*, 105(A12), 27833–27840.
- 491 Olsen, N., & Kuvshinov, A. (2004). Modeling the ocean effect of geomagnetic
492 storms. *Earth, planets and space*, 56(5), 525–530.
- 493 Pallochia, G., Amata, E., Consolini, G., Marcucci, M., & Bertello, I. (2006). Geo-
494 magnetic Dst index forecast based on IMF data only. *Annales Geophysicae*, 24(3),
495 989–999.
- 496 Patel, K., Singh, A., Singh, S. B., & Singh, A. K. (2019). Causes responsible for
497 intense and severe storms during the declining phase of Solar Cycle 24. *Journal*
498 *of Astrophysics and Astronomy*, 40(1), 1–9. Retrieved from [https://doi.org/](https://doi.org/10.1007/s12036-018-9569-7)
499 [10.1007/s12036-018-9569-7](https://doi.org/10.1007/s12036-018-9569-7) doi: 10.1007/s12036-018-9569-7
- 500 Pinkus, A. (1999). Approximation theory of the MLP model in neural networks.
501 *Acta numerica*, 8, 143–195.
- 502 Püthe, C., Manoj, C., & Kuvshinov, A. (2014). Reproducing electric field observa-
503 tions during magnetic storms by means of rigorous 3-D modelling and distortion
504 matrix co-estimation. *Earth, Planets and Space*, 66(1), 162.
- 505 Revallo, M., Valach, F., Hejda, P., & Bochníček, J. (2014). A neural network Dst
506 index model driven by input time histories of the solar wind-magnetosphere inter-
507 action. *Journal of Atmospheric and Solar-Terrestrial Physics*, 110, 9–14.
- 508 Sharifie, J., Lucas, C., & Araabi, B. N. (2006). Locally linear neurofuzzy modeling

- 509 and prediction of geomagnetic disturbances based on solar wind conditions. *Space*
510 *Weather*, 4(6).
- 511 Shen, X. C., Hudson, M. K., Jaynes, A. N., Shi, Q., Tian, A., Claudepierre,
512 S. G., ... Sun, W. J. (2017). Statistical study of the storm time radiation
513 belt evolution during Van Allen Probes era: CME- versus CIR-driven storms.
514 *Journal of Geophysical Research: Space Physics*, 122(8), 8327–8339. doi:
515 10.1002/2017JA024100
- 516 Shprits, Y. Y., Vasile, R., & Zhelavskaya, I. S. (2019). Now-casting and predicting
517 the kp index using historical values and real-time observations. *Space Weather*,
518 17. Retrieved from [https://agupubs.onlinelibrary.wiley.com/doi/abs/](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018SW002141)
519 10.1029/2018SW002141 doi: 10.1029/2018SW002141
- 520 Stepanova, M., Antonova, E., & Troshichev, O. (2005). Prediction of Dst variations
521 from polar cap indices using time-delay neural network. *Journal of atmospheric*
522 *and solar-terrestrial physics*, 67(17-18), 1658–1664.
- 523 Stepanova, M., & Pérez, P. (2000). Autoprediction of Dst index using neural net-
524 work techniques and relationship to the auroral geomagnetic indices. *GEOFISICA*
525 *INTERNACIONAL-MEXICO-*, 39(1), 143–146.
- 526 Temerin, M., & Li, X. (2006). Dst model for 1995–2002. *Journal of Geophysical Re-*
527 *search: Space Physics*, 111(A4).
- 528 Tobiska, W., Knipp, D., Burke, W., Bouwer, D., Bailey, J., Odstroil, D., ... Bow-
529 man, B. (2013). The Anemomilos prediction methodology for Dst. *Space Weather*,
530 11(9), 490–508.
- 531 Wei, H.-L., Zhu, D.-Q., Billings, S. A., & Balikhin, M. A. (2007). Forecasting
532 the geomagnetic activity of the Dst index using multiscale radial basis function
533 networks. *Advances in Space Research*, 40(12), 1863–1870.
- 534 Wolf, R., Freeman Jr, J., Hausman, B., Spiro, R., Hilmer, R., & Lambour, R.
535 (1997). Modeling convection effects in magnetic storms. *Magnetic Storms*, 98,
536 161–172.
- 537 Wu, J.-G., & Lundstedt, H. (1996). Prediction of geomagnetic storms from so-
538 lar wind data using Elman recurrent neural networks. *Geophysical research letters*,
539 23(4), 319–322.
- 540 Wu, J.-G., & Lundstedt, H. (1997). Geomagnetic storm predictions from solar wind
541 data with the use of dynamic neural networks. *Journal of Geophysical Research:*

542 *Space Physics*, 102(A7), 14255–14268.

543 Zhang, J., Richardson, I., Webb, D., Gopalswamy, N., Huttunen, E., Kasper, J., ...

544 others (2007). Solar and interplanetary sources of major geomagnetic storms (Dst

545 ≤ -100 nT) during 1996–2005. *Journal of Geophysical Research: Space Physics*,

546 112(A10).