

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. Almost all applications of neural networks to storm forecasting have utilized solar wind observations from the Earth-Sun first Lagrangian point (L1) or closer and 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: even a single layer neural network with sufficient “neurons” is ca-
 88 pable of approximating any continuous function to arbitrary precision (Leshno et al., 1993;
 89 Pinkus, 1999). However, given that this sufficient number of neurons in a single layer net-
 90 work is typically unknown and potentially intractable, workers have found success by in-
 91 stead adding layers of neurons rather than neurons themselves. This composition of lay-
 92 ers, in which neurons in a given layer operate on the outputs of neurons from the pre-
 93 ceding one, is coined “deep learning”, and has met with unprecedented success in clas-
 94 sification and regression problems. While still poorly understood beyond a heuristic sense,
 95 some workers hypothesize that deep neural networks are successful because many learn-
 96 ing problems are outcomes of hierarchical and compositional processes, which deep net-
 97 works can efficiently reproduce (Brahma et al., 2016; Lin et al., 2017). Furthermore, Lin
 98 et al. (2017) demonstrate how the properties of symmetry, locality, and polynomial log-
 99 probability in many natural processes are efficiently learned by even relatively shallow
 100 (i.e., consisting of a handful of hidden layers) neural networks.

101 **1.3 Prior applications of deep learning to geomagnetic storm forecast-** 102 **ing**

103 Previous work with NNs has focused almost entirely on prediction of Dst or other
 104 indices of geomagnetic activity, such as the Kp and the auroral electrojet (AE) indices.
 105 Supplemental Table S1 provides a succinct review of the application of NNs to the fore-

106 casting of Dst (Andriyas & Andriyas, 2015; Bala & Reiff, 2012; Gleisner et al., 1996; Jankovičová
 107 et al., 2002; Kugblenu et al., 1999; Lazzús et al., 2017; Munsami, 2000; Palocchia et al.,
 108 2006; Revallo et al., 2014; Sharifie et al., 2006; Stepanova et al., 2005; Stepanova & Pérez,
 109 2000; Wei et al., 2007; Wu & Lundstedt, 1996, 1997). These studies have applied a va-
 110 riety of architectures and data sources, but in generating forecasts for Dst, most have
 111 used the basic solar wind parameter measurements as well as prior values of Dst. All pre-
 112 vious studies applying NNs to Dst forecasting to our knowledge have utilized observa-
 113 tions made at the Earth-Sun L1 point or closer, with the exception of Chakraborty &
 114 Morley (2020), who include solar x-ray fluxes. Furthermore, almost all studies to date
 115 using NNs to forecast Dst (or any other geomagnetic storm index) have been determin-
 116 istic, generating predictions without any measure of uncertainty. Gruet et al. (2018) as-
 117 sess uncertainty in their forecasts via a Gaussian process model with fixed kernel param-
 118 eters, and this process takes as input their deterministic NN forecasts. Chakraborty &
 119 Morley (2020) on the other hand use a deep long short term memory (LSTM) network
 120 to learn how to dynamically update the kernel parameters for a Gaussian process rep-
 121 resentation of the Kp index, which is how they generate probabilistic forecasts. Finally,
 122 while not utilizing neural networks, Gu et al. (2019) generate probabilistic forecasts of
 123 the auroral electrojet (AE) index by considering output from an ensemble of 100 non-
 124 linear autoregressive models trained on independently resampled subsets of their data.

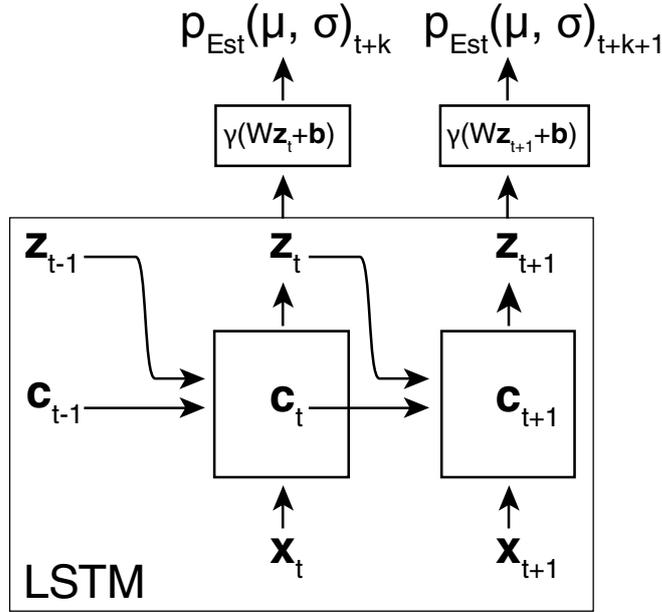
125 This work improves on previous advances by presenting the first application of prob-
 126 abilistic neural networks that explicitly generate measures of uncertainty in their out-
 127 put. Our networks are capable of learning how confident to be in their predictions, and
 128 in doing so improve forecast reliability. These networks take as input not just observa-
 129 tions from the L1 point but also observations of radiative phenomena on the solar disk.
 130 Finally, instead of forecasting Dst, we focus on its external component, Est, which does
 131 not incorporate the effects of Earth’s subsurface conductivity structure.

132 **2 Data and Methods**

133 **2.1 Probabilistic Neural Network Architecture**

134 Recently, recurrent architectures for time series regression have emerged that com-
 135 bine ridge functions with state vectors to create units with “memory”. The most suc-
 136 cessful of these has been the long short-term memory architecture (LSTM), introduced

137 by Hochreiter & Schmidhuber (1997). The LSTM cell, as its name implies, uses new in-
 138 put data with both the previous output and previous internal state to update its inter-
 139 nal state and generate new output (Supplement, Text S6). This architecture has been
 140 applied to Dst forecasting by Gruet et al. (2018), but, like all previous applications of
 141 neural networks to storm forecasting (summarized in the Supplement, Table S1), the net-
 142 work generated deterministic output with no prediction of forecast uncertainty.



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

147 We present an architecture (Figure 1) in which the NN learns to assess uncertainty
 148 in its own forecast, thereby generating probabilistic forecasts. The two basic layers uti-
 149 lized within this architecture are LSTM and dense layers. The former is described above,
 150 and the latter is an implementation of the so-called “fully connected hidden layer”, which
 151 references the fact that each entry in \mathbf{z} depends on all of the outputs from the preced-
 152 ing layer via W . That is, a dense layer that receives inputs $\mathbf{x} \in \mathbb{R}^n$ from a preceding
 153 network layer in turn generates an output vector $\mathbf{z} \in \mathbb{R}^m$ via the operation $\mathbf{z} = \gamma(W\mathbf{x} +$
 154 $\mathbf{b})$ with $W \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$, where n is the dimensionality of the preceding hid-

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

157 Inputs into our NN architecture are fed directly to an LSTM cell, and outputs from
 158 the LSTM cell are fed through a series of fully-connected hidden layers. The outputs from
 159 the last hidden layer are parameters for an output distribution over Est. We choose to
 160 use a Gaussian output distribution and compare other alternatives in the Supplement
 161 (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:

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

162 where the distribution parameters μ and σ depend on the network weights and biases
 163 and can thus be differentiated against them. However, when learning two-parameter dis-
 164 tributions, the parameter for scale often introduces leniency in the output distribution,
 165 allowing the network to expand uncertainty in its forecast rather than move its estimate
 166 for the center (see supplement Text S3).

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

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

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

171 For all implementations, training, and testing of neural networks, we use Python
 172 wrappers for the learning framework TensorFlow (Abadi et al., 2015). This framework
 173 is capable of representing neurons and the functional operations relating them as well

174 as numerically computing the relevant gradients to train the network. TensorFlow provides
 175 an implementation of the high level deep learning Keras API (<https://keras.io/>),
 176 which allows for modular construction of networks from the layers described above. We
 177 also make use of the recently introduced TensorFlow Probability library, which provides
 178 a straightforward means of adding probability distributions as layers, allowing outputs
 179 from previous layers to be used as parameters for the distribution layers. These layers
 180 are compiled into a model that contains all the operations of the entire network as well
 181 as the particular cost functions and optimizers that dictate the learning process for given
 182 training inputs and outputs. We use the Adam optimizer for gradient update steps (Kingma
 183 & Ba, 2014). All neural network configuration and training parameters are listed in Sup-
 184 plement Text S4.

185 2.2 Output Data

186 Most geomagnetic storm forecasting thus far has emphasized prediction of Dst. How-
 187 ever, Dst is actually a sum of internal and external components (Equation 3), and the
 188 internal component, Ist, is generated by currents induced in Earth’s subsurface by vari-
 189 ation in the external component, Est (Maus & Weidelt, 2004) ,

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

and

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

190 where $Q(t - \tau)$ is the impulse response that depends on subsurface electrical conduc-
 191 tivity σ , assuming that $\sigma \equiv \sigma(r)$ varies only along radial direction. This decomposi-
 192 tion is easier to express in frequency domain, in which $Q(t)$ becomes the transfer func-
 193 tion that relates internal and external components such that $\tilde{Ist}(\omega) = \tilde{Q}(\omega)\tilde{Est}(\omega)$. Sub-
 194 sequently, Est and Ist can be computed as (Maus & Weidelt, 2004):

$$\tilde{Est}(\omega) = \frac{1}{1 + \tilde{Q}(\omega)} \tilde{Dst}(\omega) \quad (5)$$

$$\tilde{Ist}(\omega) = \frac{\tilde{Q}(\omega)}{1 + \tilde{Q}(\omega)} \tilde{Dst}(\omega) \quad (6)$$

195 with knowledge of $\tilde{Q}(\omega)$ and observations of Dst. For more details about this decompo-
 196 sition and the induction transfer functions, including generalization to three-dimensional
 197 conductivity models, the reader is referred to Grayver et al. (2020); Maus & Weidelt (2004);
 198 Olsen et al. (2005).

199 The problem of forecasting Dst is then actually two separate problems: the first
 200 is forecasting Est, and the second is learning Earth’s induction response, Ist. With a suit-
 201 able model of Earth’s subsurface conductivity structure, however, knowledge of Est is
 202 sufficient to reconstruct Ist and thereby Dst. Furthermore, because the external field is
 203 what responds to magnetospheric activity anyway, it is more natural to forecast Est than
 204 Dst. Therefore, we generate forecasts of Est, and the data accessed from NOAA were
 205 generated following the methodology of Maus & Weidelt (2004) .

206 This approach will be increasingly important as we attempt to forecast higher or-
 207 der structure in Earth’s external field. Est captures only the first zonal (dipole) com-
 208 ponent of external magnetic field variability, but significant variation exists on shorter
 209 spatial scales, where the interaction with local conductivity structures becomes more im-
 210 portant and complicated (Kelbert, 2020). Given that the ultimate goal of geomagnetic
 211 storm forecasting is to forecast GICs, it is important to note that GICs themselves de-
 212 pend strongly on local conductivity structures and local external magnetic field variabil-
 213 ity (Olsen & Kuvshinov, 2004; Pütke et al., 2014). The first step to forecasting these higher
 214 order external field coefficients is forecasting a single external field coefficient, Est, which
 215 is the focus of this work.

216 **2.3 Input Data**

217 Two basic observation types relevant to geomagnetic storm forecasting have been
 218 made for the past few decades: the first includes measurements of the solar wind made
 219 in situ around the L1 point in the Earth-Sun system, and the second are measurements
 220 made directly of the solar disk and corona. These two data streams provide related but
 221 temporally disjoint information. Radiative phenomena on the solar disk take under nine
 222 minutes to be observed at Earth, while the solar wind requires two to five days to prop-
 223 agate from the solar disk to the L1 point. The L1 point is only 1.5×10^6 km from Earth,
 224 however, which is approximately one hour travel time at typical solar wind speeds (mean
 225 solar wind speed is roughly 440 km s^{-1} from the OMNI dataset).

226 Thus, while solar wind measurements near Earth’s magnetosphere are ultimately
 227 the most relevant quantities for accurate geomagnetic storm forecasting, using only ob-
 228 servations from the L1 point limits the forecast time to roughly an hour (Shprits et al.,
 229 2019). On the other hand, solar activity is ultimately responsible for all geomagnetic storms,

230 but identifying which phenomena on the solar disk have the potential to cause geomag-
 231 netic storms and predicting the storm lag times and amplitudes resulting from those phe-
 232 nomena are not trivial tasks. Observations from the solar disk include measurements of
 233 coronal mass ejections (CMEs) around the perimeter of the disk, images of the solar sur-
 234 face at various wavelengths, and surface radiative fluxes.

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

238 The low (hourly) resolution OMNI data include several solar wind and solar ob-
 239 servations, which we extracted for the years 1995-2018. During this time interval, mea-
 240 surements of the solar wind (SW) come from the WIND, IMP8, Geotail, and ACE satel-
 241 lite missions . The quantities that we use as input data from this dataset are the three
 242 components of the interplanetary magnetic field in geocentric solar magnetospheric (GSM)
 243 coordinates, SW velocity, SW particle density, SW temperature, and SW longitude and
 244 latitude incident on Earth’s magnetosphere. Because the OMNI dataset contains obser-
 245 vations from near-Earth (e.g., IMP8) and L1 (e.g., WIND, ACE) spacecraft, observa-
 246 tions are propagated to the bow shock by adding time shifts that account for the space-
 247 craft location and solar wind flow. These time shifts are included in the publicly avail-
 248 able dataset, and we utilize the observation time stamps as they are reported.

249 The GOES mission provided two time series of x-ray fluxes integrated over the so-
 250 lar disk. One series covered the wavelengths from 0.5-4 Å, and the other series covered
 251 wavelengths 1-8 Å. Given that these measurements vary over orders of magnitude, we
 252 take their logarithm as input. These data were reliably available from 1986 onwards.

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

261 We did not attempt to time shift observations of the solar disk to the bow shock
262 as is done for satellite observations at the L1 point in the OMNI dataset. This time lag
263 between the solar disk and Earth’s magnetosphere is precisely a learning problem of great
264 interest that NN’s may be able to help solve.

265 Finally, previous observations of Est were included as input while forecasting fu-
266 ture values. In total then data is available roughly from 1996 through 2018. Of these 22
267 years, we take 18 years (approximately 158,000 hours) as training data, and 4 years for
268 testing data (approximately 35,000 hours).

269 2.4 Evaluating network performance

270 In most geomagnetic storm forecasting to date, forecast accuracy is assessed by met-
271 rics such as the root mean square error and the Pearson correlation coefficient between
272 forecasted and observed data. However, these statistics are generally not compared to
273 those of a null hypothesis, for example persistence forecasting in which the best estimate
274 for any time in the future is simply the last observed value. Due to the auto-correlative
275 nature of the Est time series, persistence forecasting actually generates deceptively high
276 correlations and low errors (Figure 3, Supplement Figure S9) (Shprits et al., 2019). In
277 fact, all the networks in Supplement Table S1 either underperform or barely outperform
278 persistence forecasting as quantified by these two metrics. Furthermore, these metrics
279 are computed over both quiet and disturbed times, while the ability to predict storms
280 is the task of interest. Other than refining consideration of these metrics to only storm
281 main phases, it is not obvious what metric best evaluates forecasting performance for
282 models with deterministic output.

283 With probabilistic networks, however, reliability curves provide a useful and eas-
284 ily interpretable metric to evaluate forecast performance. Each curve corresponds to a
285 threshold in Est, and the axes compare the observed probability of exceeding that thresh-
286 old compared to the predicted probability. A perfectly reliable forecast would generate
287 curves that fall on the 1:1 line through the origin for all thresholds. If the observed re-
288 liability curve plots over the 1:1 line, then the forecast underestimates the occurrence
289 of events exceeding that threshold, while if the reliability curve plots under the 1:1 line,
290 the forecast is conservative and overestimates the occurrence of events exceeding the given
291 threshold. Because these statistics are computed for thresholds in Est, the reliability as-

292 assessment method by construction evaluates storm time forecasting separately from quiet
 293 time forecasting. We utilize reliability curves to assess forecast performance.

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

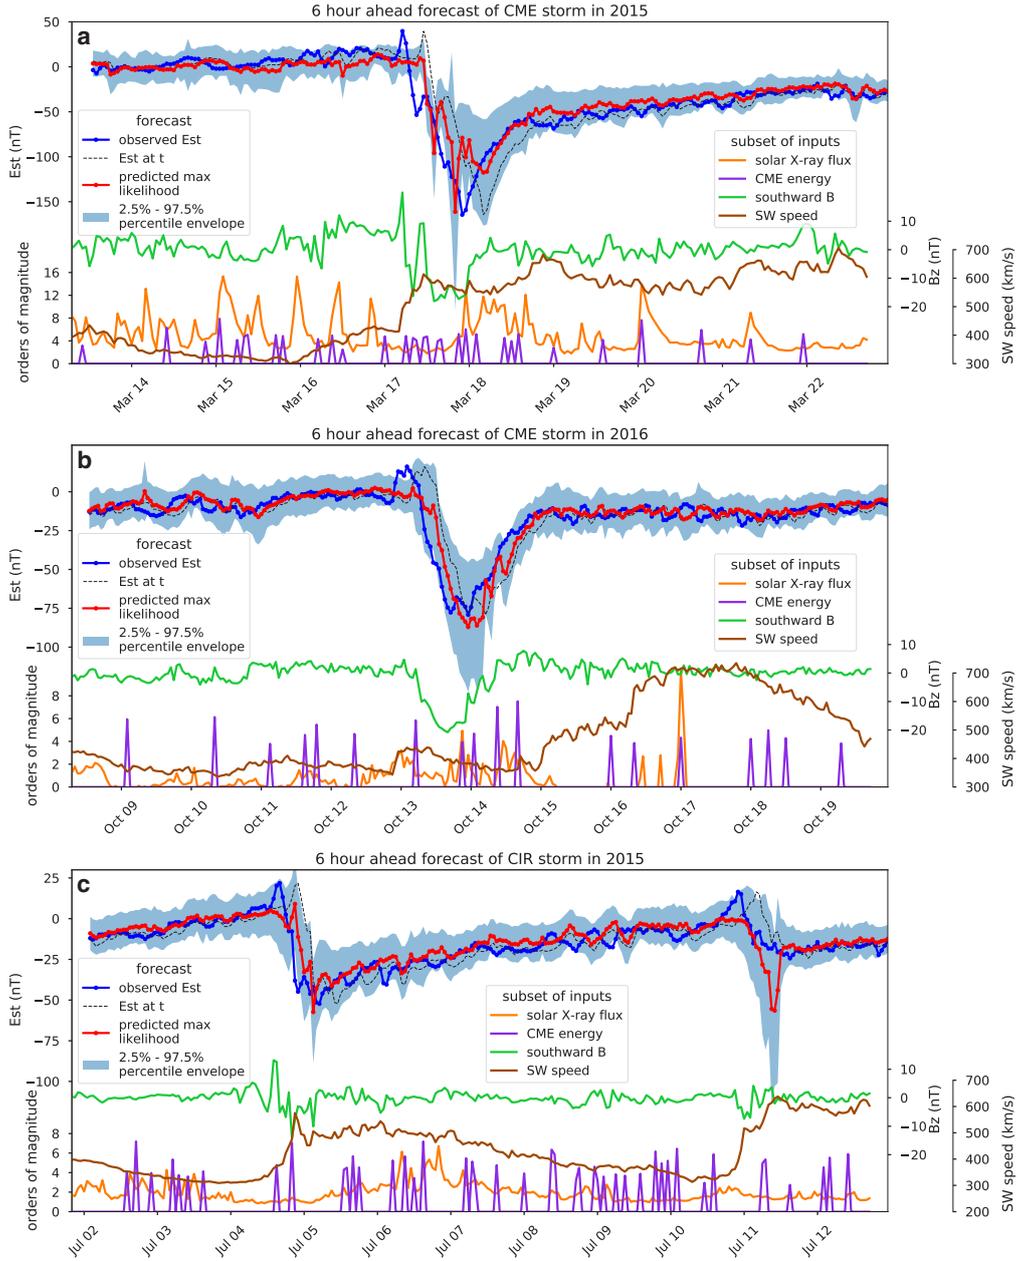
303 **3 Results and Discussion**

304 In general, our architecture is capable of learning meaningful and reliable measures
 305 of uncertainty in its forecasts, and our forecasts outperform the basic benchmark of per-
 306 sistence forecasting. We focus in our discussion on the performance of the six-hour ahead
 307 probabilistic forecasts because this forecast window is long enough that information from
 308 the L1 point is insufficient to forecast storm arrivals, allowing us to evaluate whether our
 309 network is able to leverage information from the solar disk.

310 **3.1 Storm Case Studies**

311 We present and discuss in this section six hour ahead forecasts for three different
 312 storms: two caused by CMEs (Figure 2A & B) and one resulting from a co-rotating in-
 313 teraction region (CIR) (Figure 2C). One to six hour ahead forecasts, as well as detailed
 314 comparison of the networks trained with only L1 data and both L1 and solar disk in-
 315 puts, can be found in the Supplement.

324 The first storm, in March of 2015, provides a prototypical example of a geomag-
 325 netic storm caused by a magnetic cloud emitted by the mass ejection visible as the spike
 326 in CME energy at the beginning of March 15 (Figure 2A, note that the axis is orders
 327 of magnitude) (Patel et al., 2019). The energetic mass ejection is associated with a peak
 328 in integrated x-ray fluxes, followed two days later by a relatively large geomagnetic storm
 329 beginning on March 17. The storm main phase is associated with a sustained southward



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

330 IMF of roughly -20 nT and roughly doubled solar wind speeds. In this situation, given
331 the clear connection between the mass ejection and the ensuing storm, we would expect
332 a successfully trained network to be able to expand uncertainty in its forecast as con-
333 ceivable storm arrival times approach, reflecting an understanding of the causal associ-
334 ation between activity on the solar disk and geomagnetic storms. Yet, forecast uncer-
335 tainties only expand when disturbed solar wind reaches the L1 point. At that time, the
336 network becomes aware of the storm arrival and adjusts its output by dropping Est fore-
337 casts and increasing forecast uncertainty (Figure 2A). The same is true for the smaller
338 amplitude CME storm of October 13, 2016 in Figure 2B, where forecast uncertainty only
339 grows as soon as the storm arrives at the L1 point. This storm is associated with the CME
340 visible on October 10 (Patel et al., 2019), so the occurrence of other CMEs of similar mag-
341 nitude (e.g. on October 9 and 11) demonstrates non-uniqueness that illustrates why the
342 network struggles to identify geoeffective solar activity from the provided inputs.

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

353 In all cases, the six hour ahead forecast fails to capture storm onset, during which
354 the network's forecasts tend to lag observations by the forecast length (thereby more closely
355 tracking the persistence curve) until the storm arrives at the L1 point. At that point,
356 the forecast begins to deviate from persistence as the network knows that a storm is oc-
357 ccurring. This inability to predict storm onset indicates that the network is unable to uti-
358 lize observations from the solar disk for storm arrival, which remains an open challenge.

359 Recovery is generally well-predicted, and forecasts deviate from persistence, mean-
360 ing that the network is not just taking the last observed Est value for its next forecast.
361 Unlike previous results, our network is capable of generating meaningful estimates of un-

certainty in its forecasts. In all cases, once the network detects the possible onset of a geomagnetic storm, it expands its forecast uncertainty, generally maintaining observed Est values within the 95% forecast confidence interval and providing reliable multiple hour ahead forecasts (see Supplement Text S5 for one to six hour ahead forecasts). After storm main phases, forecast uncertainties decrease during the generally well-predicted recovery phase. Given that the recovery phase is dictated by the internal dynamics of the ring current decay (and thereby independent of the solar wind state) (Daglis, 2007), its predictability is reasonable. Thus, our network exhibits forecast uncertainties that are consistent with where one would anticipate the greatest uncertainty in geomagnetic storm development with information from the L1 point, namely, the storm onset and main phase.

3.2 Conventional Metrics of Forecast Skill

In terms of the conventional forecasting skill assessments (i.e., forecast-observation Pearson correlation coefficients and root-mean-square errors, RMSE's) for one to six hour ahead forecasts, our networks outperform all previous neural network forecasts for all forecasts lengths (Figure 3). However, given that persistence forecasting for Est results in higher correlation coefficients and lower RMSE's than for persistence forecasting of Dst, our improvements should not be compared with previous results for forecasting Dst but with persistence forecasting of Est. When considering all observations, we slightly underperform persistence forecasting of Est in terms of the correlation coefficient, while outperforming in terms of RMSE, which is consistently lower. During storm times, however, our forecasting skill is much better than persistence forecasting in both metrics at all forecasting windows. The networks with both L1 and solar disk inputs always outperform networks with only L1 inputs when evaluated over both quiet and storm times (Figure S9). However, the difference in skill is small, and when considering only storm times, the five to six hour ahead forecasts of networks with only L1 inputs achieve lower RMSE's. These results again indicate that information from the solar disk does not significantly improve storm forecast skill. Finally, for particularly large storms exceeding $\text{Est} \leq -200$ nT, a forecast saturation effect is observed (Figure S9), similar to that that occurs at smaller values with different cost functions (Supplement Text S3). This effect can be partially mitigated by fine-tuning of the cost function to further facilitate the fore-

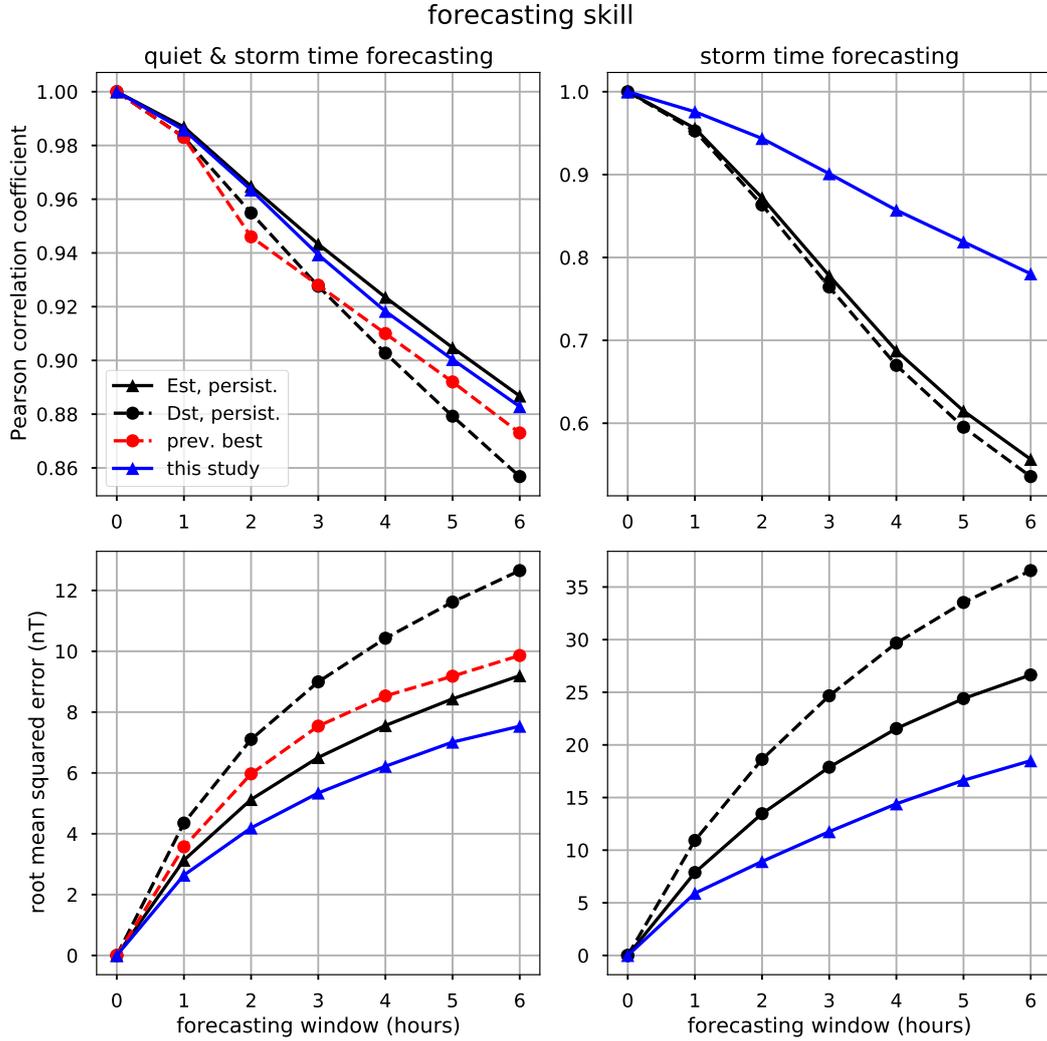
393 casting of rare storms. However, since there is only a handful of such events in the data,
394 this behaviour of the network is natural.

403 **3.3 Forecast Reliability**

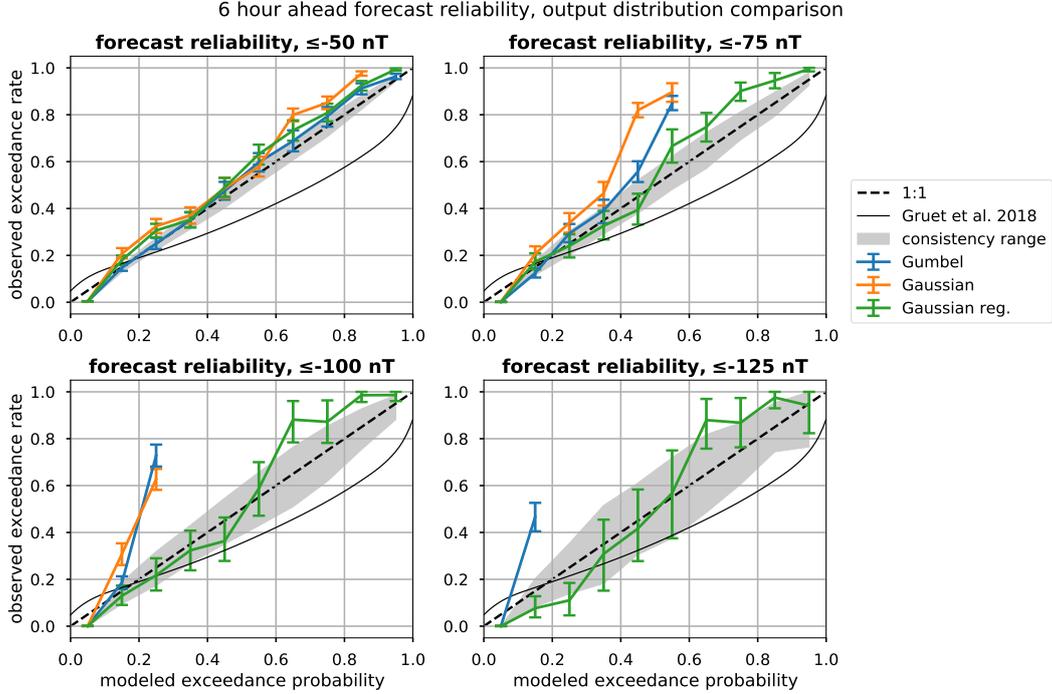
404 Reliabilities for four different storm thresholds generally overlap with the consis-
405 tency intervals for each bin, demonstrating that our network generates reliable forecasts
406 (Figure 4). For threshold of -75 and -100 nT, forecasted exceedance probabilities in the
407 range of 0.7-0.9 tended to slightly underestimate observed exceedance rates, which is con-
408 sistent with the observation that storm onsets remain difficult to predict exactly.

409 Notably, the regularization of the cost function for a Gaussian output distribution
410 significantly improves forecast reliability. Networks trained with unregularized Gaussian
411 and Gumbel output distributions (Supplement Text S3) are unable move the location
412 parameters of their forecasts during large amplitude storms, preferring instead to expand
413 forecast uncertainty, meaning that peak storm times, while often within the 95% con-
414 fidence interval, are only predicted at extremely low exceedance probabilities. This be-
415 havior explains why the reliability curves lack data to bin at high forecast probabilities
416 and furthermore why storms are underestimated for lower exceedance probabilities (Fig-
417 ure 4).

429 While some improvement in forecast reliability for smaller magnitude storms (Est
430 thresholds of -50 and -75 nT) does seem to result from the incorporation of data from
431 the solar disk (Supplement Figure S7), the preceding discussion and the result that fore-
432 cast behavior does not qualitatively change by adding solar disk inputs (Supplement Text
433 S5) indicates that we are unable to successfully utilize observations from the solar disk
434 to forecast storm arrival and amplitude. This shortcoming suggests that the informa-
435 tion necessary for identifying geoeffective solar activity is lacking in the training data,
436 and/or that the network architecture is inadequate for utilizing these data. For instance,
437 the x-ray fluxes are integrated over the entire solar disk, but peaks in these fluxes can
438 often be associated with flare events, which themselves often occur simultaneously with
439 geoeffective mass ejections (Tobiska et al., 2013). Larger, more central flares are asso-
440 ciated with larger geomagnetic storms, so adding time series of flare occurrences with
441 locations on the solar disk would complement the input series of x-ray fluxes and CMEs
442 (Tobiska et al., 2013). Furthermore, the CME dataset only includes ejections visible around



395 **Figure 3.** Conventional metrics of forecasting skill for networks with both solar disk and L1
 396 inputs (see Figure S9 for networks with only L1 inputs). The first row shows the Pearson correlation
 397 coefficient between the forecasted Est (mean, for our study) or Dst and the observed Est
 398 or Dst. The second row shows the corresponding root mean squared errors. In the columns, we
 399 show these quantities for all observations (first column) and only storm-time observations (second
 400 column). The black lines show the metrics for persistence forecasts of Est and Dst, and the red
 401 line (available only for all observations) shows the best reported performance from NN forecasts
 402 (see Table S1 for references).



418 **Figure 4.** Reliability curves for the networks with a Gumbel output and cost (blue), Gaus-
 419 sian output and cost (orange), and Gaussian output with regularized cost (green). All curves
 420 are for a six hour ahead forecast for four Est thresholds in eleven bins. Exceedance was taken in
 421 the negative sense, i.e., taking values less than or equal to the given threshold. Error bars show
 422 the 2.5-97.5 percentile range from bootstrapped resampling (number of bootstrapped samples
 423 was 1000) within the bins of forecasted exceedance probability. The envelopes show the 2.5-97.5
 424 percentile range from consistency resampling of a perfectly reliable forecast, demonstrating the
 425 conceivable range in reliable forecasts given the number of data in each bin (Bröcker & Smith,
 426 2007). The regularized Gaussian network is the most reliable of the three. Also shown is the
 427 reliability curve from Gruet et al. (2018) for their six hour ahead Dst forecast, for which the
 428 exceedance threshold is unspecified.

443 the rim of the solar disk, while geoeffective ejections occur towards the center. Thus, only
444 centralized ejections that also emit an observable lobe beyond the rim of the solar disk
445 could be reliably associated with geomagnetic storms, potentially rendering the CME
446 database largely irrelevant for the problem of geomagnetic storm forecasting. Finally,
447 integrated solar x-ray flux peaks from flares have been empirically related to solar wind
448 speeds and geomagnetic storm amplitudes, thereby providing a means of learning lag times
449 between solar activity and storm arrivals (Tobiska et al., 2013). However, LSTM net-
450 works struggle with learning lag times (Gers et al., 2002), so the network architecture
451 we have utilized is not amenable to this task.

452 **4 Conclusions**

453 This work has demonstrated a NN architecture capable of learning reliable mea-
454 sures of uncertainty in its forecasts of geomagnetic storms. Learning uncertainty in NN
455 output results in more useful probabilistic forecasts than learning uncertainty in the NN
456 parameters, and the choice of output distribution and cost function has a large impact
457 on the resulting reliability of the trained network. Specifically, adding regularizing terms
458 in the likelihood cost function improves the forecast reliability by incentivizing networks
459 to forecast more reasonable mean values rather than simply increasing forecast uncer-
460 tainty.

461 These neural networks utilize as inputs observations from both the solar disk and
462 L1 point, slightly improving forecast reliability and skill with respect to networks trained
463 only with L1 inputs. However, storm arrival and amplitude forecasting did not substan-
464 tially improve from the inclusion of these data. Thus, leveraging time series of observa-
465 tions of the solar disk, which are often sparse, remains an open problem, and future net-
466 work architectures must be carefully designed to utilize these data sources.

467 **Acknowledgments**

468 This work was partially supported by the ESA through the Swarm DISC project. The
469 low (hourly) resolution OMNI data were obtained from the GSFC/SPDF OMNIWeb in-
470 terface (<https://omniweb.gsfc.nasa.gov>). We acknowledge use of the LASCO SOHO
471 CME catalog (https://cdaw.gsfc.nasa.gov/CME_list/), which is generated and main-
472 tained at the CDAW Data Center by NASA and The Catholic University of America in
473 cooperation with the Naval Research Laboratory. SOHO is a project of international co-

474 operation between ESA and NASA. We acknowledge use of GOES mission x-ray flux data,
 475 accessed from <https://satdat.ngdc.noaa.gov/sem/goes/data/avg/>. Finally, we ac-
 476 knowledge use of the EST-IST-DST dataset from NOAA accessed from [https://www](https://www.ngdc.noaa.gov/geomag/est_ist.shtml)
 477 [.ngdc.noaa.gov/geomag/est_ist.shtml](https://www.ngdc.noaa.gov/geomag/est_ist.shtml). All data and analysis presented in this study
 478 are available as Jupyter notebooks at <https://doi.org/10.5281/zenodo.3751682>.

479 References

- 480 Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng,
 481 X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*.
 482 Retrieved from <https://www.tensorflow.org/> (Software available from tensor-
 483 flow.org)
- 484 Andriyas, T., & Andriyas, S. (2015). Relevance vector machines as a tool for fore-
 485 casting geomagnetic storms during years 1996–2007. *Journal of Atmospheric and*
 486 *Solar-Terrestrial Physics*, *125*, 10–20.
- 487 Bala, R., & Reiff, P. (2012). Improvements in short-term forecasting of geomagnetic
 488 activity. *Space Weather*, *10*(6).
- 489 Bartels, J., Heck, N. H., & Johnston, H. F. (1939). The three-hour-range index mea-
 490 suring geomagnetic activity. *Journal of Geophysical Research*, *44*(4), 411. Re-
 491 trieved from <http://doi.wiley.com/10.1029/TE044i004p00411> doi: 10.1029/
 492 TE044i004p00411
- 493 Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015). Weight uncer-
 494 tainty in neural networks. *arXiv preprint arXiv:1505.05424*.
- 495 Bolduc, L. (2002). GIC observations and studies in the hydro-Québec power
 496 system. *Journal of Atmospheric and Solar-Terrestrial Physics*, *64*(16), 1793 -
 497 1802. Retrieved from [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S1364682602001281)
 498 [S1364682602001281](http://www.sciencedirect.com/science/article/pii/S1364682602001281) (Space Weather Effects on Technological Systems) doi:
 499 [https://doi.org/10.1016/S1364-6826\(02\)00128-1](https://doi.org/10.1016/S1364-6826(02)00128-1)
- 500 Brahma, P. P., Wu, D., & She, Y. (2016). Why Deep Learning Works: A Manifold
 501 Disentanglement Perspective. *IEEE Transactions on Neural Networks and Learn-*
 502 *ing Systems*, *27*(10), 1997–2008. doi: 10.1109/TNNLS.2015.2496947
- 503 Bröcker, J., & Smith, L. A. (2007, jun). Increasing the Reliability of Reliability
 504 Diagrams. *Weather and Forecasting*, *22*(3), 651–661. Retrieved from [http://](http://journals.ametsoc.org/doi/abs/10.1175/WAF993.1)
 505 journals.ametsoc.org/doi/abs/10.1175/WAF993.1 doi: 10.1175/WAF993.1

- 506 Burton, R. K., McPherron, R., & Russell, C. (1975). An empirical relationship be-
 507 tween interplanetary conditions and Dst. *Journal of geophysical research*, *80*(31),
 508 4204–4214.
- 509 Chakraborty, S., & Morley, S. K. (2020, jul). Probabilistic prediction of geomag-
 510 netic storms and the K $\{p\}$ index. *Journal of Space Weather and Space Cli-*
 511 *mate*, *10*, 36. Retrieved from <http://arxiv.org/abs/2007.02733>[https://](https://www.swsc-journal.org/10.1051/swsc/2020037)
 512 www.swsc-journal.org/10.1051/swsc/2020037 doi: 10.1051/swsc/2020037
- 513 Daglis, I. A. (2007, feb). Ring Current Dynamics. *Space Science Reviews*, *124*(1-4),
 514 183–202. Retrieved from [http://link.springer.com/10.1007/s11214-006-9104-](http://link.springer.com/10.1007/s11214-006-9104-z)
 515 [z](http://link.springer.com/10.1007/s11214-006-9104-z) doi: 10.1007/s11214-006-9104-z
- 516 Daglis, I. A., Thorne, R. M., Baumjohann, W., & Orsini, S. (1999). The terres-
 517 trial ring current: Origin, formation, and decay. *Reviews of Geophysics*, *37*(4),
 518 407–438.
- 519 Gers, F. A., Schraudolph, N. N., & Schmidhuber, J. (2002). Learning precise tim-
 520 ing with LSTM recurrent networks. *Journal of Machine Learning Research*, *3*(1),
 521 115–143. doi: 10.1162/153244303768966139
- 522 Gleisner, H., Lundstedt, H., & Wintoft, P. (1996). Predicting geomagnetic storms
 523 from solar-wind data using time-delay neural networks. *Annales Geophysicae*, *14*,
 524 679.
- 525 Gonzalez, W. D., Tsurutani, B. T., & De Gonzalez, A. L. C. (1999). Interplanetary
 526 origin of geomagnetic storms. *Space Science Reviews*, *88*(3-4), 529-562.
- 527 Grayver, A. V., Kuvshinov, A. V., & Werthmüller, D. (2020). Time-domain mod-
 528 elling of 3-D Earth’s and planetary EM induction effect in ground and satellite
 529 observations. *in prep.*
- 530 Gruet, M., Chandorkar, M., Sicard, A., & Camporeale, E. (2018). Multiple-hour-
 531 ahead forecast of the Dst index using a combination of long short-term memory
 532 neural network and Gaussian process. *Space Weather*, *16*(11), 1882–1896.
- 533 Gu, Y., Wei, H. L., Boynton, R. J., Walker, S. N., & Balikhin, M. A. (2019). Sys-
 534 tem Identification and Data-Driven Forecasting of AE Index and Prediction Un-
 535 certainty Analysis Using a New Cloud-NARX Model. *Journal of Geophysical*
 536 *Research: Space Physics*, *124*(1), 248–263. doi: 10.1029/2018JA025957
- 537 Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural compu-*
 538 *tation*, *9*(8), 1735–1780.

- 539 Jankovičová, D., Dolinský, P., Valach, F., & Vörös, Z. (2002). Neural network-
540 based nonlinear prediction of magnetic storms. *Journal of atmospheric and solar-*
541 *terrestrial physics*, *64*(5-6), 651–656.
- 542 Kelbert, A. (2020). *The Role of Global/Regional Earth Conductivity Models in*
543 *Natural Geomagnetic Hazard Mitigation* (Vol. 41) (No. 1). Springer Netherlands.
544 Retrieved from <https://doi.org/10.1007/s10712-019-09579-z> doi: 10.1007/
545 s10712-019-09579-z
- 546 Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*
547 *preprint arXiv:1412.6980*.
- 548 Kugblenu, S., Taguchi, S., & Okuzawa, T. (1999). Prediction of the geomagnetic
549 storm associated Dst index using an artificial neural network algorithm. *Earth*
550 *Planets Space*, *51*(307-313).
- 551 Lazzús, J., Vega, P., Rojas, P., & Salfate, I. (2017). Forecasting the Dst index using
552 a swarm-optimized neural network. *Space Weather*, *15*(8), 1068–1089.
- 553 Leshno, M., Lin, V. Y., Pinkus, A., & Schocken, S. (1993). Multilayer feed-
554 forward networks with a nonpolynomial activation function can approximate
555 any function. *Neural Networks*, *6*(6), 861 - 867. Retrieved from [http://](http://www.sciencedirect.com/science/article/pii/S0893608005801315)
556 www.sciencedirect.com/science/article/pii/S0893608005801315 doi:
557 [https://doi.org/10.1016/S0893-6080\(05\)80131-5](https://doi.org/10.1016/S0893-6080(05)80131-5)
- 558 Lin, H. W., Tegmark, M., & Rolnick, D. (2017, sep). Why Does Deep and Cheap
559 Learning Work So Well? *Journal of Statistical Physics*, *168*(6), 1223–1247. Re-
560 trieved from <http://link.springer.com/10.1007/s10955-017-1836-5> doi: 10
561 .1007/s10955-017-1836-5
- 562 Love, J. J., Lucas, G. M., Kelbert, A., & Bedrosian, P. A. (2018). Geoelectric
563 Hazard Maps for the Mid-Atlantic United States: 100 Year Extreme Values
564 and the 1989 Magnetic Storm. *Geophysical Research Letters*, *45*(1), 5–14. doi:
565 10.1002/2017GL076042
- 566 Maus, S., & Weidelt, P. (2004). Separating the magnetospheric disturbance mag-
567 netic field into external and transient internal contributions using a 1D con-
568 ductivity model of the Earth. *Geophysical Research Letters*, *31*(12), 2–5. doi:
569 10.1029/2004GL020232
- 570 Munsami, V. (2000). Determination of the effects of substorms on the storm-
571 time ring current using neural networks. *Journal of Geophysical Research: Space*

- 572 *Physics*, 105(A12), 27833–27840.
- 573 Olsen, N., & Kuvshinov, A. (2004). Modeling the ocean effect of geomagnetic
574 storms. *Earth, planets and space*, 56(5), 525–530.
- 575 Olsen, N., Sabaka, T. J., & Lowes, F. (2005). New parameterization of external and
576 induced fields in geomagnetic field modeling, and a candidate model for igrf 2005.
577 *Earth, planets and space*, 57(12), 1141–1149.
- 578 Pallochia, G., Amata, E., Consolini, G., Marcucci, M., & Bertello, I. (2006). Geo-
579 magnetic Dst index forecast based on IMF data only. *Annales Geophysicae*, 24(3),
580 989–999.
- 581 Patel, K., Singh, A., Singh, S. B., & Singh, A. K. (2019). Causes responsible for
582 intense and severe storms during the declining phase of Solar Cycle 24. *Journal*
583 *of Astrophysics and Astronomy*, 40(1), 1–9. Retrieved from [https://doi.org/](https://doi.org/10.1007/s12036-018-9569-7)
584 10.1007/s12036-018-9569-7 doi: 10.1007/s12036-018-9569-7
- 585 Pinkus, A. (1999). Approximation theory of the MLP model in neural networks.
586 *Acta numerica*, 8, 143–195.
- 587 Pütthe, C., Manoj, C., & Kuvshinov, A. (2014). Reproducing electric field observa-
588 tions during magnetic storms by means of rigorous 3-D modelling and distortion
589 matrix co-estimation. *Earth, Planets and Space*, 66(1), 162.
- 590 Revallo, M., Valach, F., Hejda, P., & Bochníček, J. (2014). A neural network Dst
591 index model driven by input time histories of the solar wind-magnetosphere inter-
592 action. *Journal of Atmospheric and Solar-Terrestrial Physics*, 110, 9–14.
- 593 Sharifie, J., Lucas, C., & Araabi, B. N. (2006). Locally linear neurofuzzy modeling
594 and prediction of geomagnetic disturbances based on solar wind conditions. *Space*
595 *Weather*, 4(6).
- 596 Shen, X. C., Hudson, M. K., Jaynes, A. N., Shi, Q., Tian, A., Claudepierre,
597 S. G., . . . Sun, W. J. (2017). Statistical study of the storm time radiation
598 belt evolution during Van Allen Probes era: CME- versus CIR-driven storms.
599 *Journal of Geophysical Research: Space Physics*, 122(8), 8327–8339. doi:
600 10.1002/2017JA024100
- 601 Shprits, Y. Y., Vasile, R., & Zhelavskaya, I. S. (2019). Now-casting and predicting
602 the kp index using historical values and real-time observations. *Space Weather*,
603 17. Retrieved from [https://agupubs.onlinelibrary.wiley.com/doi/abs/](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018SW002141)
604 10.1029/2018SW002141 doi: 10.1029/2018SW002141

- 605 Stepanova, M., Antonova, E., & Troshichev, O. (2005). Prediction of Dst variations
 606 from polar cap indices using time-delay neural network. *Journal of atmospheric
 607 and solar-terrestrial physics*, *67*(17-18), 1658–1664.
- 608 Stepanova, M., & Pérez, P. (2000). Autoprediction of Dst index using neural net-
 609 work techniques and relationship to the auroral geomagnetic indices. *GEOFISICA
 610 INTERNACIONAL-MEXICO-*, *39*(1), 143–146.
- 611 Temerin, M., & Li, X. (2006). Dst model for 1995–2002. *Journal of Geophysical Re-
 612 search: Space Physics*, *111*(A4).
- 613 Tobiska, W., Knipp, D., Burke, W., Bouwer, D., Bailey, J., Odstrcil, D., . . . Bow-
 614 man, B. (2013). The Anemomilos prediction methodology for Dst. *Space Weather*,
 615 *11*(9), 490–508.
- 616 Wei, H.-L., Zhu, D.-Q., Billings, S. A., & Balikhin, M. A. (2007). Forecasting
 617 the geomagnetic activity of the Dst index using multiscale radial basis function
 618 networks. *Advances in Space Research*, *40*(12), 1863–1870.
- 619 Wolf, R., Freeman Jr, J., Hausman, B., Spiro, R., Hilmer, R., & Lambour, R.
 620 (1997). Modeling convection effects in magnetic storms. *Magnetic Storms*, *98*,
 621 161–172.
- 622 Wu, J.-G., & Lundstedt, H. (1996). Prediction of geomagnetic storms from so-
 623 lar wind data using Elman recurrent neural networks. *Geophysical research letters*,
 624 *23*(4), 319–322.
- 625 Wu, J.-G., & Lundstedt, H. (1997). Geomagnetic storm predictions from solar wind
 626 data with the use of dynamic neural networks. *Journal of Geophysical Research:
 627 Space Physics*, *102*(A7), 14255–14268.
- 628 Zhang, J., Richardson, I., Webb, D., Gopalswamy, N., Huttunen, E., Kasper, J., . . .
 629 others (2007). Solar and interplanetary sources of major geomagnetic storms (Dst
 630 ≤ -100 nT) during 1996–2005. *Journal of Geophysical Research: Space Physics*,
 631 *112*(A10).