

Forecasting GICs and geoelectric fields from solar wind data using LSTMs: application in Austria

R. L. Bailey¹, R. Leonhardt¹, C. Möstl², C. Beggan³, M. A. Reiss²,
A. Bhaskar⁴, A. J. Weiss^{2,5}

¹Conrad observatory, Zentralanstalt für Meteorologie und Geodynamik, Vienna, Austria

²Space Research Institute, Austrian Academy of Sciences, Graz, Austria

³British Geological Survey, Edinburgh, UK

⁴Space Physics Laboratory, ISRO/Vikram Sarabhai Space Centre, Trivandrum, India

⁵Institute of Physics, University of Graz, Universitätsplatz 5, 8010 Graz, Austria

Key Points:

- The aim is to directly forecast GICs rather than dB/dt , which is often used as a proxy.
- Results from LSTMs predicting either E_x and E_y or substation GICs from solar wind data are compared.
- GIC forecasting seems to work best when the LSTM model is trained directly on GIC data.

Corresponding author: Rachel Bailey, r.bailey@zang.ac.at

Abstract

The forecasting of local GIC effects has largely relied on the forecasting of dB/dt as a proxy and, to date, little attention has been paid to directly forecasting the geoelectric field or GICs themselves. We approach this problem with machine learning tools, specifically recurrent neural networks or LSTMs by taking solar wind observations as input and training the models to predict two different kinds of output: first, the geoelectric field components E_X and E_Y ; and second, the GICs in specific substations in Austria. The training is carried out on the geoelectric field and GICs modelled from 26 years of one-minute geomagnetic field measurements, and results are compared to GIC measurements from recent years. The GICs are generally predicted better by an LSTM trained on values from a specific substation, but only a fraction of the largest GICs are correctly predicted. This model had a correlation with measurements of around 0.6, and a root-mean-square error of 0.7 A. The probability of detecting mild activity in GICs is around 50%, and 15% for larger GICs.

Plain Language Summary

Using satellites, we measure the state of the solar wind a short distance away from the Earth (at the so-called Lagrange-1 or L1 point) to see what is coming towards us at any given moment. Changes in the solar wind such as an increase in wind speed or a strong magnetic field can potentially impact satellite operation in orbit and power grid infrastructure on the ground - in extreme cases, solar storms can damage power grids and transformers by inducing electrical currents in the power lines. These are called geomagnetically induced currents (GICs). Here, we attempt to forecast the scales of GICs by applying machine learning methods, specifically Long-Short-Term-Memory recurrent neural networks, to take the solar wind data measured at the L1 point and predict the currents that would be seen in power grids in Austria. This gives us a lead time of around 10 to 40 minutes in the forecast. We discuss whether it is best to attempt to predict the ground electric field that leads to the GICs or the GICs themselves, and discuss the difficulties in this kind of prediction and the shortfalls in the model.

1 Introduction

(Bailey et al., n.d.)

Geomagnetically induced currents (GICs) have long been known to affect power grids, transformers and any earthed conductive networks spanning large distances (for an overview, see Boteler et al., 1998; Boteler & Pirjola, 2017; Kelbert, 2020). GICs can cause problems in power grid operation such as transformer overheating or permanent transformer damage and system collapse in extreme cases (Molinski, 2002), leading to further societal and economic harm (Eastwood et al., 2018). Although studies of GICs were restricted to high latitudes where the consequences are more pronounced, mid-latitudes are being paid increasingly more attention as local effects such as transformer overheating are discovered (Barbosa et al., 2015; Butala et al., 2017; Lotz & Danskin, 2017; Gil et al., 2019; Caraballo et al., 2020; Svanda, Michal et al., 2020, among others).

The forecasting of GICs has developed alongside studies into the effects of regional GICs (Pulkkinen et al., 2006). Forecasting in particular is a complex problem due to the chain of cascading induction effects from the impingement of solar wind at the bow shock down to currents flowing between the earth and power grids on the surface. Improving predictive GIC modelling is listed as one of the open questions still to address to achieve GIC readiness (Pulkkinen et al., 2017).

Most studies so far have focused on predicting geomagnetic activity - such as dB/dt , which is often used as a proxy for GICs - from solar wind data measured at L1 or in near-

65 Earth space. The earliest studies addressing this problem with neural network architec-
 66 ture are Wintoft (2005) and Wintoft et al. (2015), followed by Lotz and Cilliers (2015)
 67 and recently Keesee et al. (2020) and Tasistro-Hart et al. (2021). The *Dst/SYMH* in-
 68 dex in particular has received a lot of attention from geophysicists and machine learn-
 69 ing engineers alike (e.g. Lu et al., 2016; Bhaskar & Vichare, 2019; Wintoft & Wik, 2021).

70 While dB/dt is often used as a proxy for GICs, it does not provide the whole pic-
 71 ture. The downside of modelling with this approach is that dB/dt only functions as a
 72 useful indicator of GIC activity. The relationship between dB/dt and E (which is the
 73 primary factor determining the scale of the GICs) depends on the magnetotelluric trans-
 74 fer function, which is frequency dependent (Chave & Jones, 2012). Single values of the
 75 time derivative of the magnetic field can only be useful GIC proxies if further assump-
 76 tions on the frequency content are made (Pulkkinen et al., 2006).

77 What do we do if we want to develop a model that provides forecasts that power
 78 grid operators can work with? One approach would be to directly forecast the surface
 79 geoelectric field, from which GICs at different stations can be calculated. In compari-
 80 son to the many studies into forecasting dB/dt and *Dst*, little effort has been devoted
 81 to forecasting geoelectric fields thus far. Pulkkinen et al. (2009, 2010) studied the fore-
 82 casting of GICs from remote solar observations, allowing a few days warning before larger
 83 events. Modelling of geoelectric fields from solar wind to ground using full MHD mod-
 84 elling has been carried out by Pulkkinen, Hesse, et al. (2007), Zhang et al. (2015) and
 85 Honkonen et al. (2018), and with empirical modelling in Lotz et al. (2017).

86 In this study, we aim to tackle this problem from another angle and forecast re-
 87 gional GICs from L1 solar wind data using a machine learning method, and we compare
 88 the results to observations of GICs in Austria. We try this with two different approaches:
 89 in the first, we train a model to forecast the geoelectric field and calculate the GICs from
 90 there, and in the second we forecast the GICs directly. Predictions from both methods
 91 are evaluated and compared using data from recent years.

92 This study is structured as follows. Section 2 describes the data used in this study,
 93 including an analysis of 26 years of geomagnetic measurements used to model GICs in
 94 the region of Austria and a case study looking at the 2003 Halloween storm. Section 3
 95 then goes on to describe the models built to forecast GIC values, and the results are pre-
 96 sented in Section 4, discussed in Section 5 and summarised in Section 6.

97 2 Data

98 This analysis relies on INTERMAGNET-quality geomagnetic observatory data, which
 99 ensures a high quality of data with few data gaps or disturbances. We use data with a
 100 cadence of one minute because these are available for a long time period (26 years), which
 101 is not possible with 1 Hz data. Data with 1-minute resolution should be representative
 102 of most important GIC content (Pulkkinen et al., 2006). Due to Austria’s small size (roughly
 103 280 x 600 km), we assume that the geomagnetic variations are roughly constant across
 104 it both latitudinally and longitudinally, and therefore only select and use geomagnetic
 105 variations from one station at a time.

106 In the following, we describe the data sets used in this study. Geomagnetic field
 107 variations from observatory measurements were used to calculate the ground geoelectric
 108 field in Austria. GICs at any power grid substation can be calculated from the geoelec-
 109 tric field, and the equations for two specific substations are determined using a linear fit
 110 to observed GICs. In terms of the geomagnetic and geoelectric field components, x and
 111 y refer to the geographic northward and eastward directions respectively.

112 2.1 Geomagnetic observatory data from WIC and FUR

113 The Conrad Observatory (WIC), situated at a geomagnetic latitude of 42.95° and
 114 longitude of 89.94° according to AACGM-v2 (Shepherd, 2014), is located southwest of
 115 Vienna near the town of Muggendorf in Lower Austria. High quality geomagnetic mea-
 116 surements have been carried out here since the official opening mid-2014, providing six
 117 years of data for analysis. We extend the time range using data from Fürstenfeldbruck
 118 (FUR) in Bavaria, Germany. Initial studies are done using WIC data, and studies of long-
 119 term measurements are carried out using FUR data. A map showing the location of the
 120 two stations can be found in **Fig. 1**.

121 The Fürstenfeldbruck Geomagnetic Observatory (geomagnetic lat: 43.06° , lon: 85.93°)
 122 is one of the closest INTERMAGNET-quality geomagnetic observatories to the Conrad
 123 Observatory. It is situated almost directly west of WIC and separated by 348 km. This
 124 station is a very good proxy for geomagnetic field variations in Austria due to its prox-
 125 imity and the similar geomagnetic latitude and geological setting. Measurements at a
 126 quality high enough for this analysis have been carried out since 1995, providing twenty-
 127 six years of data or 13.7 million data points at a 1-minute resolution.

128 An analysis of the coherence between WIC and FUR data has been carried out for
 129 the overlapping years of measurements (2015-2021), in which the Pearson’s correlation
 130 coefficient (PCC) between the two time series doesn’t drop below 0.99 for either the x
 131 or y variables over all six years. The correlation in variations (dBx/dt and DBy/dt) is
 132 slightly lower, with the lowest values (0.91) seen in the DBy/dt values.

133 2.2 Goelectric field

134 In order to model the expected levels of GICs, we need knowledge of the ground
 135 goelectric field in the region. The goelectric field for the past 26 years is modelled di-
 136 rectly from the 1-minute geomagnetic field variations at FUR. The model approach used
 137 is the one-dimensional plane wave method (e.g. Boteler & Pirjola, 2017) using the EU-
 138 RHOM model number 39 (Ádám et al., 2012) to describe the one-dimensional layers of
 139 resistivity going into the Earth. We assume the time series is representative across the
 140 country, which is a reasonable approach for small areas but not for larger countries. The
 141 plane wave approach was used in favour of the thin-sheet approach used in previous stud-
 142 ies (Bailey et al., 2017, 2018) for the shorter computation times with similar levels of ac-
 143 curacy. The calculation results in the horizontal goelectric field components E_x and E_y .
 144 Note that the x -component in the goelectric field corresponds to the y -component ge-
 145 omagnetic field variations, and vice versa.

146 2.3 Geomagnetically induced currents

147 To evaluate the levels of GICs over the 26 years of available FUR data, we do not
 148 follow the standard modelling procedure of putting the goelectric field components through
 149 the full power grid network, which would be computationally heavy, but instead find a
 150 direct linear fit of the goelectric field components to measurements of GICs to find the
 151 current at station j , i.e.

$$GIC_j = a_j \cdot E_x + b_j \cdot E_y \quad (1)$$

152 where a_j and b_j are station-specific real coefficients (with units A·km/V). This approach
 153 can only be used on transformer stations with measurements since the coefficients must
 154 be determined from a linear fit to the data, but it often has similar or better accuracy
 155 than results from a network model. See Pulkkinen, Pirjola, and Viljanen (2007) or Torta
 156 et al. (2012) for more discussion on this method and for the equations determining a_j
 157 and b_j .

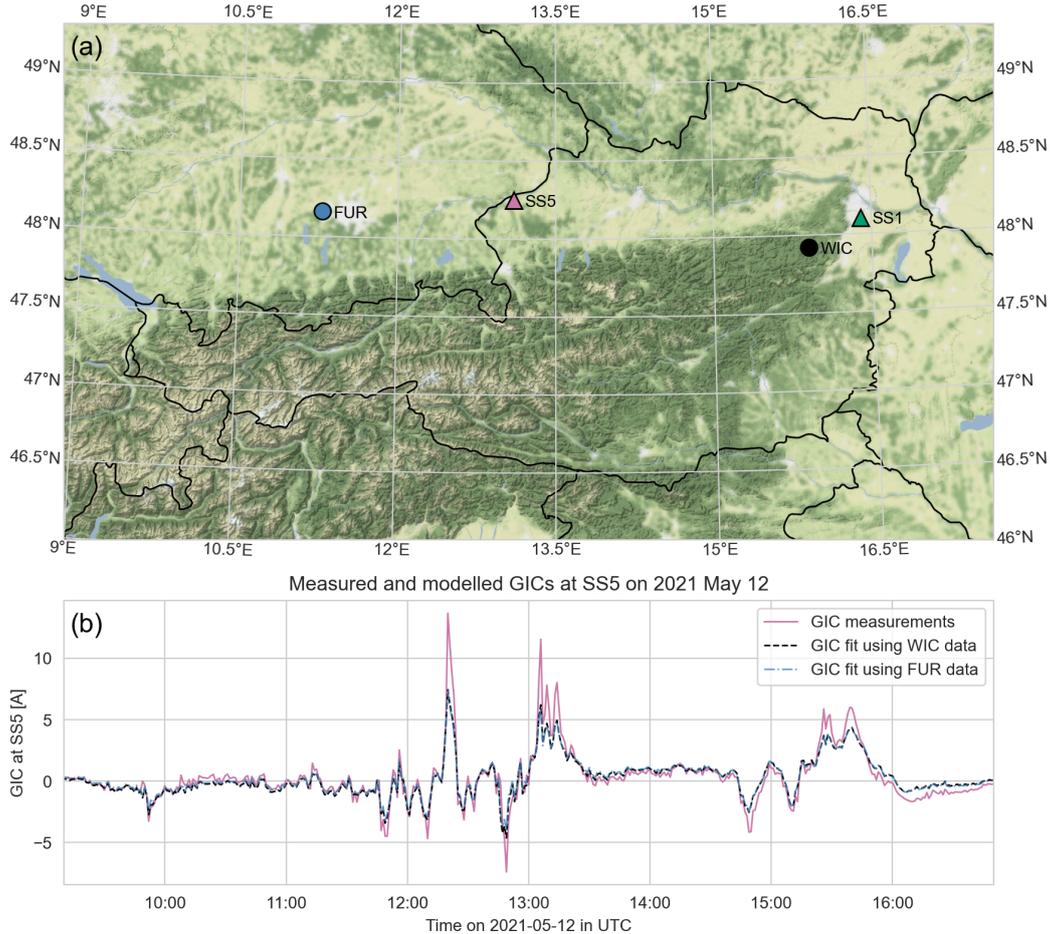


Figure 1. (a) A map showing the locations of two power grid substations (triangles) and the two geophysical observatories (circles) used for geoelectric field modelling, and (b) an example of GIC fit from modelled geoelectric field values for a geomagnetic storm in May 2021. The solid line (purple) shows transformer neutral point current measurements that have been offset-corrected and resampled via interpolation to a 1-minute sampling rate (from 1-second). The two dashed lines show the GICs calculated from E using WIC (black) and FUR (blue) data, which are nearly identical. Note that the largest GIC values are almost always underestimated despite the otherwise good agreement between model and measurements.

158 The fit for Eq. 1 was applied to measurements of direct currents from multiple trans-
 159 former neutral points in Austrian power grid substations provided by the Graz Univer-
 160 sity of Technology, a summary of which can be found in Albert et al. (2022). In this study,
 161 only measurements from two substations were used: one near Vienna (hereafter referred
 162 to SS1 for Substation 1) and another north of Salzburg (SS5), both with sampling rates
 163 of one second. The data was resampled to a one minute sampling rate for use in this study
 164 using a 1-minute median sliding window. These two stations are of interest because they
 165 are in the high-voltage network and experience larger GICs than the other stations with
 166 measurements. As such they are useful examples for depicting the expected maximum
 167 scales of GICs that could be seen across the grid. We choose three geomagnetically ac-
 168 tive periods and use the geoelectric field components E_x and E_y modelled from FUR data
 169 to derive the following equations:

$$GIC_{SS1} = 3.77 \cdot 10^{-2} \cdot E_x + 3.19 \cdot 10^{-2} \cdot E_y \quad (2)$$

$$GIC_{SS5} = 0.44 \cdot 10^{-2} \cdot E_x + 5.55 \cdot 10^{-2} \cdot E_y \quad (3)$$

170 We see that the x -component of the geoelectric field contributes roughly the same
 171 amount to the GICs seen in SS1 as the y component. The y -component of the geoelec-
 172 tric field mostly dominates the currents in SS5 and contributes ten times more than the
 173 x -component. The differences in contributions from geoelectric field components stem
 174 from the varying grid layout and connections at each substation. An analysis shows that
 175 the GICs calculated from these equations are slightly more accurate than those from the
 176 full network model. Comparing to measurements at SS1, the Pearson's correlation co-
 177 efficients for both GICs from the network model and GICs from Eq. 1 are 0.86, while at
 178 SS5 the correlation improves from 0.85 to 0.88. In both cases the amplitudes of the GICs
 179 are better matched and the root-mean-square-errors drop from 0.24 to 0.12 A at SS1 and
 180 0.46 to 0.12 A at SS5. These measures were calculated from a fit of the geoelectric field
 181 data to measurements using eight days of geomagnetically active periods (including the
 182 September 2017 storm). This includes the most recent active period, meaning the mea-
 183 surements should represent the current grid configuration and we exclude fitting only
 184 to grid noise by using a geomagnetically active period. A fit applied to the geoelectric
 185 field modelled from WIC rather than FUR data produces slightly different coefficients
 186 but results in the same level of accuracy when compared to GIC measurements. An ex-
 187 ample of the measurements and GIC fits can be seen in **Fig. 1b**.

188 Regardless of which time range the fit is applied to, the GICs calculated using Eq.
 189 1 (as well as those from the network model) tend to underestimate the peaks of the largest
 190 GICs by up to a factor of two (see e.g. **Fig 1b**, 12:20 or 13:05 UTC). We assume this
 191 is a result of attenuation of the modelled geoelectric field due to the lower sampling rate
 192 used for field modelling (Grawe et al., 2018) or the oversimplification of using a uniform
 193 geoelectric field and 1D model of the subsurface resistivity (Ngwira et al., 2015; Sun &
 194 Balch, 2019; Weigel, 2017). Despite this, the very good agreement between model and
 195 measurements means that any results based on the modelled geoelectric fields will still
 196 be reasonable.

197 In addition to the absolute GIC values, we also look at the cumulative absolute GICs
 198 over an hour, GIC_{sum1h} . GIC_{sum1h} is taken as the sum of values over the hour divided
 199 by the number of timesteps in an hour (60 for our minute values) to make it indepen-
 200 dent of sampling rate, and is used as a separate indicator for geomagnetic activity, more
 201 representative of sustained GICs than large spikes, both of which can have different (but
 202 similarly detrimental) effects on transformers (Bolduc, 2002; Gaunt & Coetzee, 2007).
 203 Using the accumulated sum of GICs or geoelectric field has seen usage in other studies,
 204 although not often - Lotz and Danskin (2017) used the accumulated E over varying pe-
 205 riods and Viljanen et al. (2014) also worked with daily GIC sum averaged across nodes.
 206 In Austria, 0 to 0.5 Ah can be seen during quiet times, and values above that generally
 207 represent more active times.

208 2.4 Distribution of values

209 In order to determine how best to forecast GICs, we first look at the 26 years of
 210 available data and the distributions of both geomagnetic variations and modelled GICs.
 211 **Figure 2a** presents the distribution of FUR minute dBx/dt and dB_y/dt variations. There
 212 are very few values populating the tail of the distribution where the largest values are
 213 found. High values for this region are at 80 nT/min and upwards. The largest variations
 214 occur most commonly in the x -direction (leading to larger E_y) rather than the y -direction,

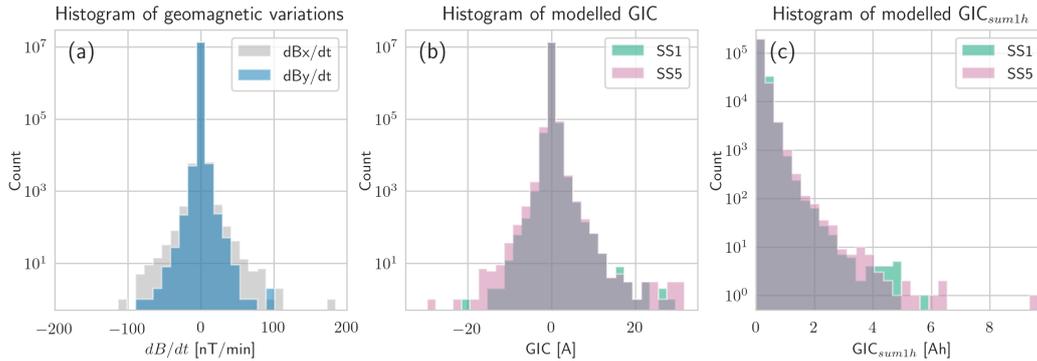


Figure 2. Histograms showing the distribution of the values in (a) the geomagnetic variations at FUR, (b) the GICs modelled from dB/dt at two substations, and (c) the hourly cumulative modelled GICs at two substations for all data, GIC_{sum1h} . The y-axes have logarithmic scales.

215 implying that stations in the power grid sitting on east-west lines are already more sus-
 216 ceptible to larger GICs.

217 In **Figs. 2b and 2c**, the GICs observed at SS5 are larger than those at SS1. While
 218 the size of the currents depends largely on the network topology and grounding resis-
 219 tance, we noted in Section 2.3 that the currents at SS5 are mostly determined by the y -
 220 component of the geoelectric field (or x -component of the geomagnetic field variations),
 221 which generally sees larger variations.

222 2.5 Most active days

223 In **Table 1**, the 10 most active days in the 26 years of data according to different
 224 measures of activity dBx/dt and dB/dt at FUR, modelled $|GIC|$ and GIC_{sum1h} at both
 225 SS1 and SS5 are listed. There are many overlapping days between the different measures,
 226 making a total of 19 days. Bold font highlights the ten largest values in each column.

227 A similar table for largest GIC days in Central Europe was produced in Viljanen
 228 et al. (2014, Table 4), and we see that the tables are very much in agreement with 17
 229 shared dates, even though the table in Viljanen et al. (2014) is only based on one vari-
 230 able. They used a value akin to the GIC_{sum1h} used here, namely the daily sum of GICs
 231 averaged across all nodes. Similarly, 17 of the days listed here also appear in Juusola et
 232 al. (2015), Table 3, where an analysis of the days with largest GICs was carried out for
 233 Northern Europe. Other larger storms that have occurred since those studies (March 2015
 234 and September 2017) do not stand out in comparison to those from the last solar cycle
 235 with the exception of the storm from June 2015.

236 The largest values in each measure are clearly centered around the 2003 Halloween
 237 storm. Large values in dBx/dt tend to go alongside large GIC values in SS5, and days
 238 with large GIC_{sum1h} usually coincide with days with larger $|GIC|$, as expected. Some
 239 exceptions are 2000-09-17, 2001-04-08, 2005-01-07 and 2005-08-24, which only show high
 240 cumulative GICs but do not stand out in dB/dt -values and peak GICs. A comparison
 241 of these events shows they have large and unidirectional geomagnetic field variations (with
 242 total field changes of 100 to 300 nT) that occur over an hour or more. These in partic-
 243 ular lead to sustained GICs in stations susceptible to geomagnetic field changes in that
 244 direction. The variations on 2000-09-17 are shown as an example of this kind of behaviour
 245 in **Fig. 3**. Although not extremely geomagnetically active, they show that power grid

Table 1. Table showing the ten most active days according to the maximum values in three measures: leftmost are the horizontal geomagnetic field variations (dBx/dt and dBy/dt), in the centre the absolute GICs ($|GIC|$) at two different transformer stations (SS1 and SS5), and rightmost the cumulative GICs over an hour at two transformer stations (GIC_{sum1h}). Bold font highlights the ten largest values seen in that measure. The largest values are seen during the Halloween Storm on 2003 October 29-31 (italicised).

Date	dBx/dt [nT/min]	dBy/dt [nT/min]	$ GIC1 $ [A]	$ GIC5 $ [A]	$GIC1_{sum1h}$ [Ah]	$GIC5_{sum1h}$ [Ah]
1998-05-04	52.0	46.0	11.37	9.56	2.48	2.81
2000-04-06	42.9	43.7	8.78	11.34	3.00	3.43
2000-07-15	184.7	28.5	20.30	29.47	4.79	6.25
2000-09-17	34.5	19.9	10.45	9.89	4.21	4.21
2001-03-31	82.4	40.7	10.85	17.55	3.69	3.18
2001-11-06	85.1	38.1	12.73	13.72	3.95	5.24
2001-11-24	62.4	33.3	14.20	17.81	4.51	4.18
<i>2003-10-29</i>	102.9	92.3	28.57	31.67	5.77	9.66
<i>2003-10-30</i>	33.1	40.3	17.68	16.44	4.66	4.71
<i>2003-10-31</i>	91.5	56.2	14.75	16.88	2.41	4.03
2003-11-20	19.8	31.4	11.63	10.73	4.82	4.62
2004-07-26	78.5	8.5	10.15	15.33	1.20	1.53
2004-11-07	43.0	37.7	7.33	8.60	2.54	2.67
2004-11-08	24.7	28.9	9.77	9.42	4.17	3.57
2004-11-09	76.1	49.9	14.21	13.70	4.28	3.46
2005-05-15	36.3	35.1	11.45	13.96	4.38	6.31
2005-08-24	41.6	31.9	10.51	13.18	4.01	6.16
2005-09-11	60.7	30.7	8.81	12.50	1.24	1.65
2015-06-22	63.0	12.8	9.94	16.67	2.56	3.47

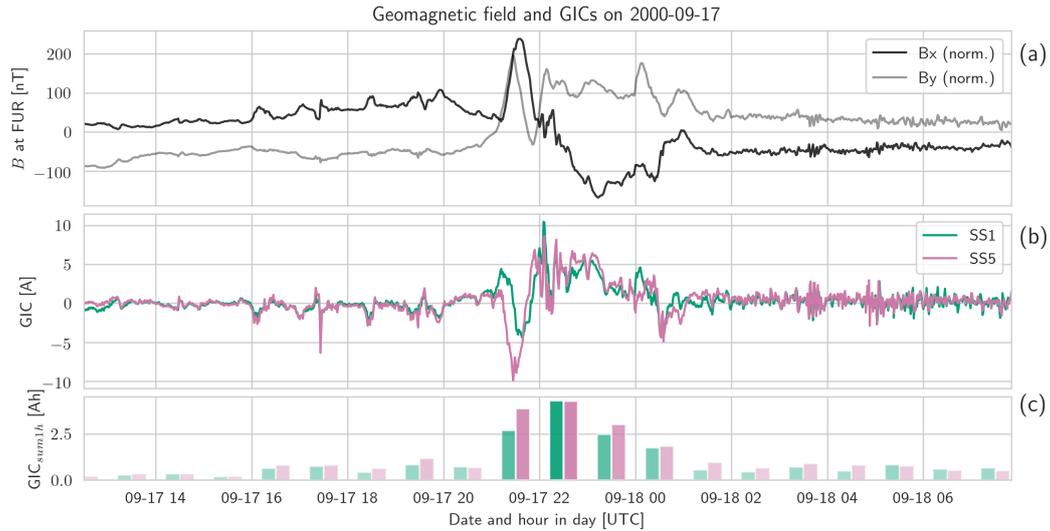


Figure 3. Plot of (a) geomagnetic variations at FUR (normalised to around zero by subtracting the mean field strength), (b) modelled GICs at two substations, and (c) cumulative hourly GICs on 2000-09-17 as an example of a day with no extreme GIC values but large cumulative hourly GICs.

246 transformers would have been subjected to large amounts of cumulative GICs sustained
 247 over an hour at least.

248 2.6 Case study: 2003 Halloween Storm

249 In **Fig. 2**, almost all of the values in the tail end of the distribution resulted from
 250 the “Halloween storm”, which lasted from 2003 October 29 to November 1. These also
 251 make up the largest values in **Table 1**, with maximum GIC values almost twice as large
 252 as the other values seen. We now conduct a detailed analysis of the behaviour during
 253 this storm and the GICs that were likely present in the power grid as an example of the
 254 problems that can arise when using only dB/dt as a proxy for GICs. We see that both
 255 large instantaneous GICs and sustained GICs appear without large dB/dt values.

256 The geomagnetic storm that occurred at the end of October in 2003 was the re-
 257 sult of a series of fast and geoeffective coronal mass ejections hitting the Earth during
 258 a particularly active period around the maximum of solar cycle 23 (e.g., Gopalswamy
 259 et al., 2005). In Eastwood et al. (2018), this storm was classified as a 1-in-10 year event,
 260 and is not considered an exceptionally rare example. No event of this or a higher mag-
 261 nitude has occurred since 2003 (with the exception of a CME directed away from Earth
 262 on July 2012, see Ngwira et al., 2013; Baker et al., 2013; Liu et al., 2014), and such events
 263 are somewhat more probable during the solar maxima (Owens et al., 2021), but have also
 264 occurred at any point throughout the solar cycle.

265 A brief evaluation of this storm for Austria was carried out in Bailey et al. (2018),
 266 in which a maximum GIC of 14 A was modelled. Using an updated model with newer
 267 data allows us to get a more accurate estimate of GICs during stronger events, and us-
 268 ing the method from Section 2.3 for SS1 and SS5 we see the values reaching 25–30 A.
 269 Taking into account that the GIC peaks modelled using minute data generally under-
 270 estimate the observations, these could also have reached up to 60 A.

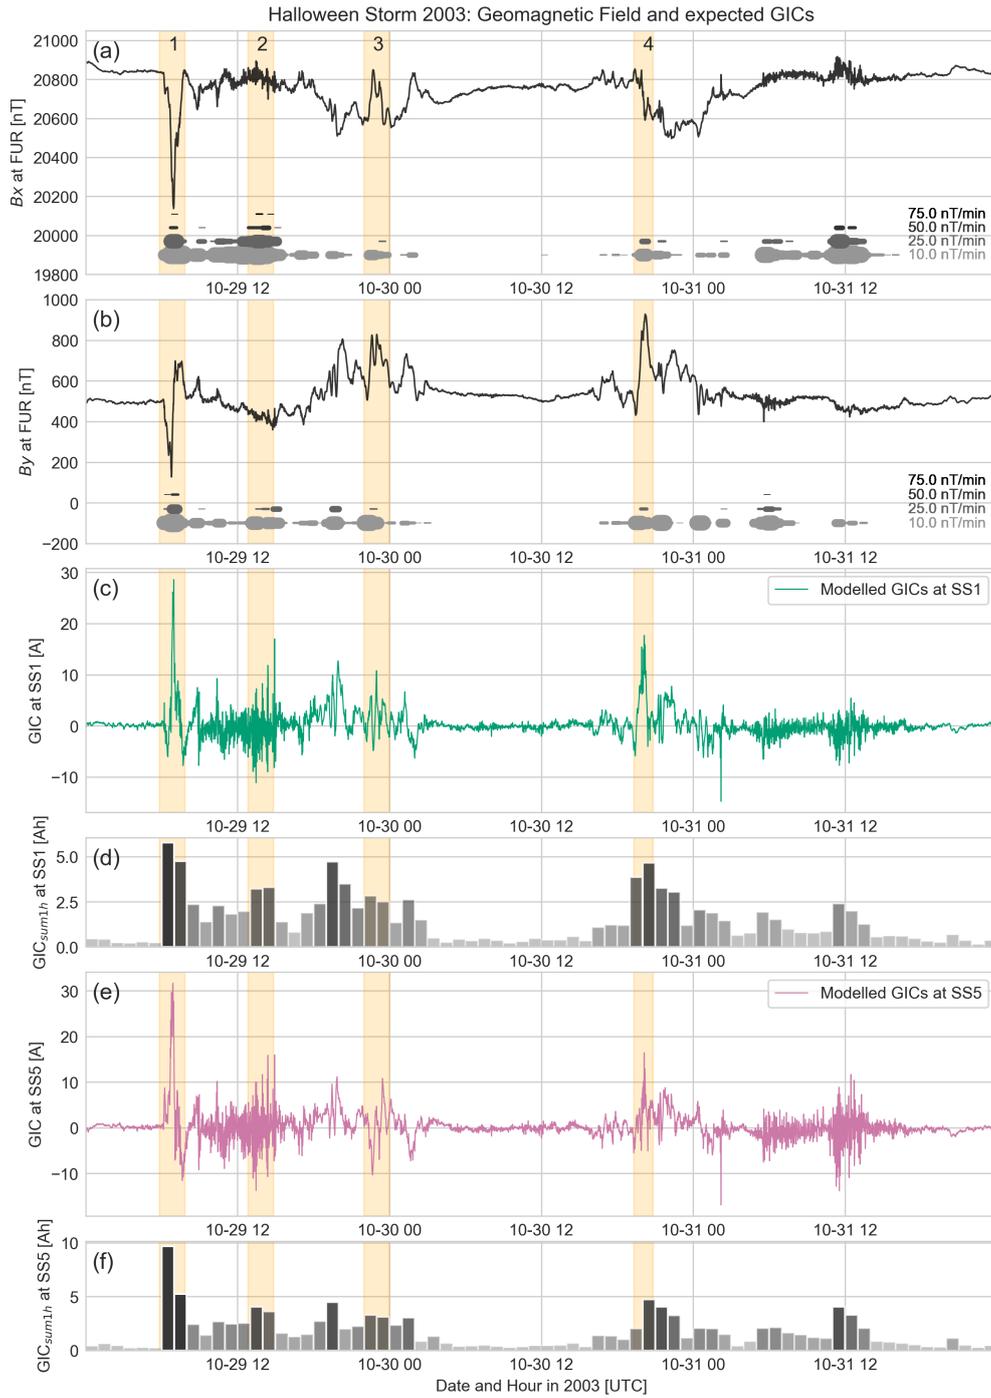


Figure 4. The Halloween storm from 2003 October 29 till 2003 November 1, during which some of the largest geomagnetic variations of the last few decades were seen. (a) and (b) show the geomagnetic variations at FUR in the x and y directions. Plotted below are levels of activity (10, 25, 50, and 75 nT/min) with line thickness showing how often these values were exceeded over a certain time range. (c) and (e) show the modelled GICs at the substations SS1 and SS5, and (d) and (f) show the cumulative GICs over each hour at each substation.

271 **Figure 4** compares the geomagnetic field and the modelled GICs for the 2003 Hal-
 272 loween Storm. Panels (a) and (b) show the geomagnetic field variations in the x and y
 273 directions. The thick lines plotted below the field show the presence of various levels of
 274 dB/dt variations (as they might be shown using a forecasting method). Light grey shows
 275 a level of 10 nT/min, and this increases going upwards to 25 nT/min, 50 nT/min and
 276 75 nT/min. The thickness of the line shows how often the value was exceeded within a
 277 time frame of 30 minutes (with a maximum being 30 times). Panels (c) and (e) show the
 278 GICs calculated from the modelled geoelectric field at the substations SS1 and SS5, and
 279 the panels (d) and (f) show the cumulative sum of absolute GIC values (GIC_{sum1h}) over
 280 1-hour periods.

281 Four time intervals, highlighted in yellow on the plot, have been picked out for dis-
 282 cussion. Intervals 1 and 2 have been selected because, as can be seen in the high levels
 283 of dB/dt in both components, these were the most active periods. Intervals 3 and 4, in
 284 contrast, were chosen because of continuously low levels of dB/dt but lack of higher ($>$
 285 50 nT/min) values.

286 Interval 1 shows a large GIC value, which is fairly short-lived. Interval 2, in con-
 287 trast, shows a consistent level of moderate GICs, though it does not reach an extremely
 288 high value. Interval 3 has a similar level of sustained GIC_{sum1h} as Interval 2 despite it
 289 having a comparatively smaller amount of dB/dt over the same period. In Interval 4,
 290 SS1 experiences the second highest value of GIC (17 A) throughout the whole storm, even
 291 though there is only continuous low-level dBx/dt and dB_y/dt (10 to 25 nT/min), most
 292 of it unidirectional (comparable to the type of signal seen in **Fig. 3**). On top of that,
 293 the cumulative GICs are also some of the highest.

294 In summary, we see there are large differences between periods that have short-lived
 295 but large GICs (Intervals 1 and 4) and those that have longer periods of sustained GICs
 296 (Intervals 2 and 3), and both large GICs and sustained GICs can appear without large
 297 dB/dt because the ground geoelectric field responds at a range of frequencies not cap-
 298 tured by dB/dt intensity alone. Each scenario could lead to different problems if it were
 299 to occur in a transformer to any large degree (Price, 2002; Gaunt & Coetzee, 2007; Bolduc,
 300 2002).

301 3 Building a Forecasting Model

302 From the analysis of past data, we deduce that, in order to forecast a comprehen-
 303 sive summary of expected GIC behaviour, we need to forecast either both geoelectric field
 304 components or the GICs directly. While the magnitude of the field is most important,
 305 the direction also plays an important role. From Eqs. 2 and 3, we see that a large value
 306 in E_x at SS5, for example, could be cancelled out by a smaller negative one in the E_y
 307 value, and the opposite could be true elsewhere, making a station-by-station approach
 308 advantageous.

309 We now move on to build a forecasting model based on these conclusions. Three
 310 machine learning methods were put through an initial comparison for evaluation: a stan-
 311 dard feed-forward neural network (NN) with three layers (32 neurons initially), a gra-
 312 dient boosting regressor based on XGBoost in Python (with 400 decision trees), and a
 313 recurrent neural network (specifically, a Long-Short-Term Memory RNN or LSTM) with
 314 three layers (32 blocks initially) and a basic Attention mechanism. The three types or
 315 architecture were set up in size and hyperparameter choice to be somewhat compar-
 316 able in basic accuracy on an initial subset of the training data set, then were provided the
 317 full, identical data sets (scaled and shaped according to each method) and compared ac-
 318 cording to a set of metrics for model evaluation (root-mean-square error, Pearson's cor-
 319 relation coefficient, probability of detection). From these first comparisons, the LSTM
 320 with Attention showed the most promise and was developed into the final model, although

321 due to the myriad machine learning methods available these days there may well be other
 322 approaches equally suited for this task. Details on the comparison can be found in the
 323 Jupyter Notebook #4 listed in Sec. 7.

324 3.1 Data preparation

325 The input to the machine learning model is solar wind data measured at L1 and
 326 forward-propagated to the bow shock. This means that, assuming we take measurements
 327 from satellites situated at L1, we have a varying forecast lead time between 15 and 60
 328 minutes depending on the solar wind speed. The high resolution OMNI data set (see sec-
 329 tion on Data Availability for details) was used for solar wind measurements (speed, den-
 330 sity, and magnetic field components) at a minute cadence combined with the local time
 331 and day in year to make up the features, while the model target was either the geoelec-
 332 tric field (E) modelled from FUR data or the GICs modelled from the E_x and E_y com-
 333 ponents.

334 Taking solar wind measurements that have already been propagated forward to the
 335 bow shock, we use the two hours prior to the time we wish to forecast as input. This goes
 336 from $t - 120$ minutes to $t - 0$, where t is the forecast time. The range of 120 minutes
 337 for past data was decided on through experimentation, where the period was increased
 338 until longer periods did not lead to any improvements in the forecasting skill. To reduce
 339 the size and complexity of the input data, it is subsampled to a 10-minute resolution by
 340 picking every 10th point (rather than interpolation and/or fitting, which we found led
 341 to a loss in forecast skill), resulting in sequences of length 12. These sequences are used
 342 as input to forecast the maximum value of E or GICs over 40 minutes from $t - 10$ to
 343 $t + 30$. This step of ten minutes into the “past” (which reduces the lead time by ten min-
 344 utes) is to account for possible timing errors in propagating the solar wind forward to
 345 the bow shock.

346 Sampling the modelled geoelectric field or GIC data to produce a balanced data
 347 set for model training is challenging because there is a clear bias towards quiet times and
 348 not enough data from geomagnetically active times (with a factor of roughly $10^7 : 1$
 349 for quiet to active). An initial approach using the entire data set led to a trained model
 350 that predicted only quiet times, which could not be remedied without additional data
 351 handling or large changes to the training methods. The target data set was therefore se-
 352 lectively sampled to reduce the imbalance. The distribution of samples was undersam-
 353 pled in the range of $E = 0$ to 100 mV/km ($GIC = 0$ to 8 A). Above that, we applied
 354 some data augmentation by duplicating the samples by 2 to 5 times and applying a ran-
 355 dom offset in time to the input data of each to avoid identical samples. The offset was
 356 randomly sampled without replacement from values between -10 and +10 minutes, which
 357 shifts the input solar wind data that the model sees, and means that the maximum value
 358 was either closer to the start or the end of the following 40-minute forecast window. Oth-
 359 erwise, all samples had a minimum time difference of 60 minutes between them. The re-
 360 sulting distribution is close to a one-sided Gaussian distribution. Roughly the same num-
 361 ber of samples (9000) were used in training for each target.

362 The samples were split into training and testing sets by time. The years 2000, 2001
 363 were reserved for validation to aid in model selection during training, while 2017, 2019
 364 and 2020 were reserved for testing, and the remaining 21 years were used in training. The
 365 presence of data gaps longer than 15 consecutive minutes in the OMNI data set led to
 366 samples being excluded from the analysis - this led to 8 to 15% sample exclusion, depend-
 367 ing on the years used. Data gaps shorter than 15 minutes were linearly interpolated over.

368 We reduced all values of $E > 200$ mV/km ($GIC > 15$ A) to 200 mV/km (15 A)
 369 because the larger values were only present in roughly 100 of the 13.7 million time-steps
 370 (or five to seven events in the 25-year period) and heavily skewed the distribution, in which
 371 all values were scaled between 0 and 1. Rescaling points above this limit greatly improved

372 the level to which the model could learn the problem but also means that the maximum
 373 forecast the model can realistically produce is for 200 mV/km. This was tested by eval-
 374 uating a model trained on data clipped at 200 mV/km versus one trained on the orig-
 375 inal data, and the model trained on clipped data performed better on both clipped and
 376 unclipped test data sets.

377 3.2 Training the LSTM

378 To approach this forecasting problem, we use a four-layer LSTM with an Atten-
 379 tion layer. The Attention mechanism is meant to simulate human attention (first devel-
 380 oped in Bahdanau et al., 2015), which can be understood intuitively as a mechanism that
 381 picks out the most important part of a sequence and discards the parts that are consid-
 382 ered irrelevant. It is a tool now commonly applied in natural language processing for ex-
 383 ample (Galassi et al., 2020). The model is structured so that the input first goes through
 384 an LSTM layer and then through the Attention mechanism. The data is then fed into
 385 another LSTM layer before going through a final feed-forward layer to reduce the out-
 386 put to a single value.

387 For geoelectric field prediction, the LSTM branches into two: the left side deals with
 388 a regression problem, namely forecasting the maximum magnitude of the geoelectric field.
 389 We chose a custom loss function for the regression problem where events (peaks) are rare
 390 in the data, and where the scale of the peaks is important. A min-max scaling factor used
 391 as a penalty term meant that training to match the peak value would drive the loss down.
 392 The right side of the LSTM forecasts the sign of the geoelectric field in a classification
 393 problem, which in this case is the sign of the maximum field value used for the regres-
 394 sion problem. Here, the binary cross-entropy loss function was used. Training worked
 395 better when the two were trained as separate targets, rather than attempting to fore-
 396 cast E without taking the absolute value first. The regression problem appears to be not
 397 too difficult a task, but the model had far more problems trying to forecast the direc-
 398 tion. In training, the weights of the two problems are, when scaled, about 15 : 1 for re-
 399 gression to classification. The classification problem to determine the sign is given sec-
 400 ondary importance because even an LSTM dedicated to this problem had trouble achiev-
 401 ing a good level of accuracy. A diagram of the different LSTM architectures, the loss func-
 402 tions and the hyperparameters used for the training of each model can be found in the
 403 supporting information. Iteration through the various possible hyperparameters was car-
 404 ried out for all four models for optimisation. Similar sets of hyperparameters were found
 405 for each LSTM application, with some minor differences between them, although the choice
 406 of the same hyperparameters for all applications also led to reasonable models in all cases.
 407 Regularisation was applied in the form of dropout.

408 Multiple models were trained to evaluate the best approach for forecasting GICs.
 409 Those trained to forecast the geoelectric field components are referred to as LSTM-E,
 410 while nets trained to forecast the GICs directly are referred to as LSTM-GIC. Both neu-
 411 ral nets are only trained on the output of geophysical models (in the case of E, the re-
 412 sult of FUR variations put through the plane-wave model, and for GICs, these are the
 413 currents calculated in power grid transformers from E) because we don't have measure-
 414 ments of E or GIC over long enough periods and because, as described in Sec. 2.3, GICs
 415 from geophysical models reach a good enough accuracy to be a reasonable substitute in
 416 training. Both models predict the absolute value of the target, but the LSTM-E predicts
 417 the sign (positive or negative) in addition.

418 3.3 Evaluating the model skill

419 Each model was trained on its respective training set and the best LSTM param-
 420 eters were chosen based on model behaviour when presented with the validation set. Fol-
 421 lowing training, we ran the model on the test data set in a virtual 'real-time mode' pro-

422 viding updates to the input data every 15 minutes, and giving an output with a 15-minute
 423 cadence. The comparison to the ground truth (either the modelled geoelectric field or
 424 measured GICs) is performed point-to-point as well as by looking at events, where the
 425 event-based analysis is given the most importance. In order to have a benchmark for com-
 426 parison, we produced a real-time persistence approach which takes the maximum of the
 427 geoelectric field or GICs in the 20 minutes before the solar wind measurement time to
 428 forecast the maximum when the solar wind would reach Earth. As such, the persistence
 429 model (PERS) also uses a varying forecast lead time. The machine-learning forecast model
 430 should be able to beat persistence in most measures.

431 Our event-based analysis follows the recommendations put forward by Pulkkinen
 432 et al. (2013) and Welling et al. (2018) for dB/dt forecasting. An “event” in the data is
 433 classified as a value that exceeds a certain threshold, while all values below that thresh-
 434 old are non-events. By defining a threshold, we can calculate the confusion matrix (Wilks,
 435 2011), which includes the number of correctly-predicted events or true positives (TP),
 436 missed events or false negatives (FN), incorrectly-predicted events or false positives (FP),
 437 and the correctly-predicted non-events or true negatives (TN). The metrics proposed in
 438 Pulkkinen et al. (2013) include the Probability of Detection (POD), which is the frac-
 439 tion of measured events correctly predicted as events, also called the true positive rate
 440 (TPR or $TP/(TP+FN)$). Similarly, we include the probability of False Detection (POFD),
 441 the fraction of measured non-events incorrectly predicted as events, which is equivalent
 442 to the false positive rate (FPR or $FP/(FP+TN)$). In addition, the Heidke Skill Score (HSS)
 443 and True Skill Statistic (TSS) are also considered, both of which are derived from all vari-
 444 ables in the confusion matrix (see e.g. Heidke, 1926; Bloomfield et al., 2012). Both the
 445 HSS and TSS show no model skill at 0, and better model skill when approaching 1. The
 446 TSS has the benefit over the HSS of being unbiased by event/non-event ratios. We also
 447 include the bias (BS), which shows if the model tends to over-predict (more false pos-
 448 itives, $BS > 1$) or under-predict (more false negatives, $BS < 1$).

449 4 Results

450 We present the results split in two parts: in the first part, we test our model’s fore-
 451 casting ability with regards to the the geoelectric field components. The results are com-
 452 pared to the geoelectric field modelled from geomagnetic variations at FUR (see Sec. 2.2).
 453 In the second part, we test the forecasting ability for GICs. These are calculated using
 454 (1) the geoelectric field components predicted from LSTM-E to calculate the GICs at
 455 the two substations we picked for analysis, and (2) directly from LSTM-GIC for each sub-
 456 station. The comparison between the model results and measurements of GICs is car-
 457 ried out for the years 2017, 2019 and 2020.

458 For the evaluation of geoelectric field forecast, we compute the scores for three event
 459 thresholds: these are 30, 60, and 90 mV/km in both E_x and E_y . In GICs, the level of
 460 60 mV/km corresponds to a current of roughly 4 A through either SS1 or SS5, and we
 461 use similar thresholds of 2, 4 and 6 A. It is difficult to determine the minimum level of
 462 GICs above which transformers may experience adverse effects because these are heav-
 463 ily dependent on transformer type and the presence of DC-handling mechanisms. We
 464 have too few measurements of GICs exceeding higher levels such as 10 A to make an anal-
 465 ysis at this level useful, but 4 A is crossed often during geomagnetically active times. The
 466 results are described in the next section.

467 **Figure 5** gives a graphical representation of the model behaviour at each thresh-
 468 old using receiver-operator characteristic (ROC) and detection-error tradeoff (DET) curves.
 469 Both depict the model’s ability to forecast events at varying thresholds. The ROC curve
 470 shows the trade-off between the true positive rate (also POD) and false positive rate (also
 471 POFD) at different event thresholds. Usually, when the threshold is low, the TPR is high
 472 but we also see an increased FPR, which is unwanted - a model that captures the ob-

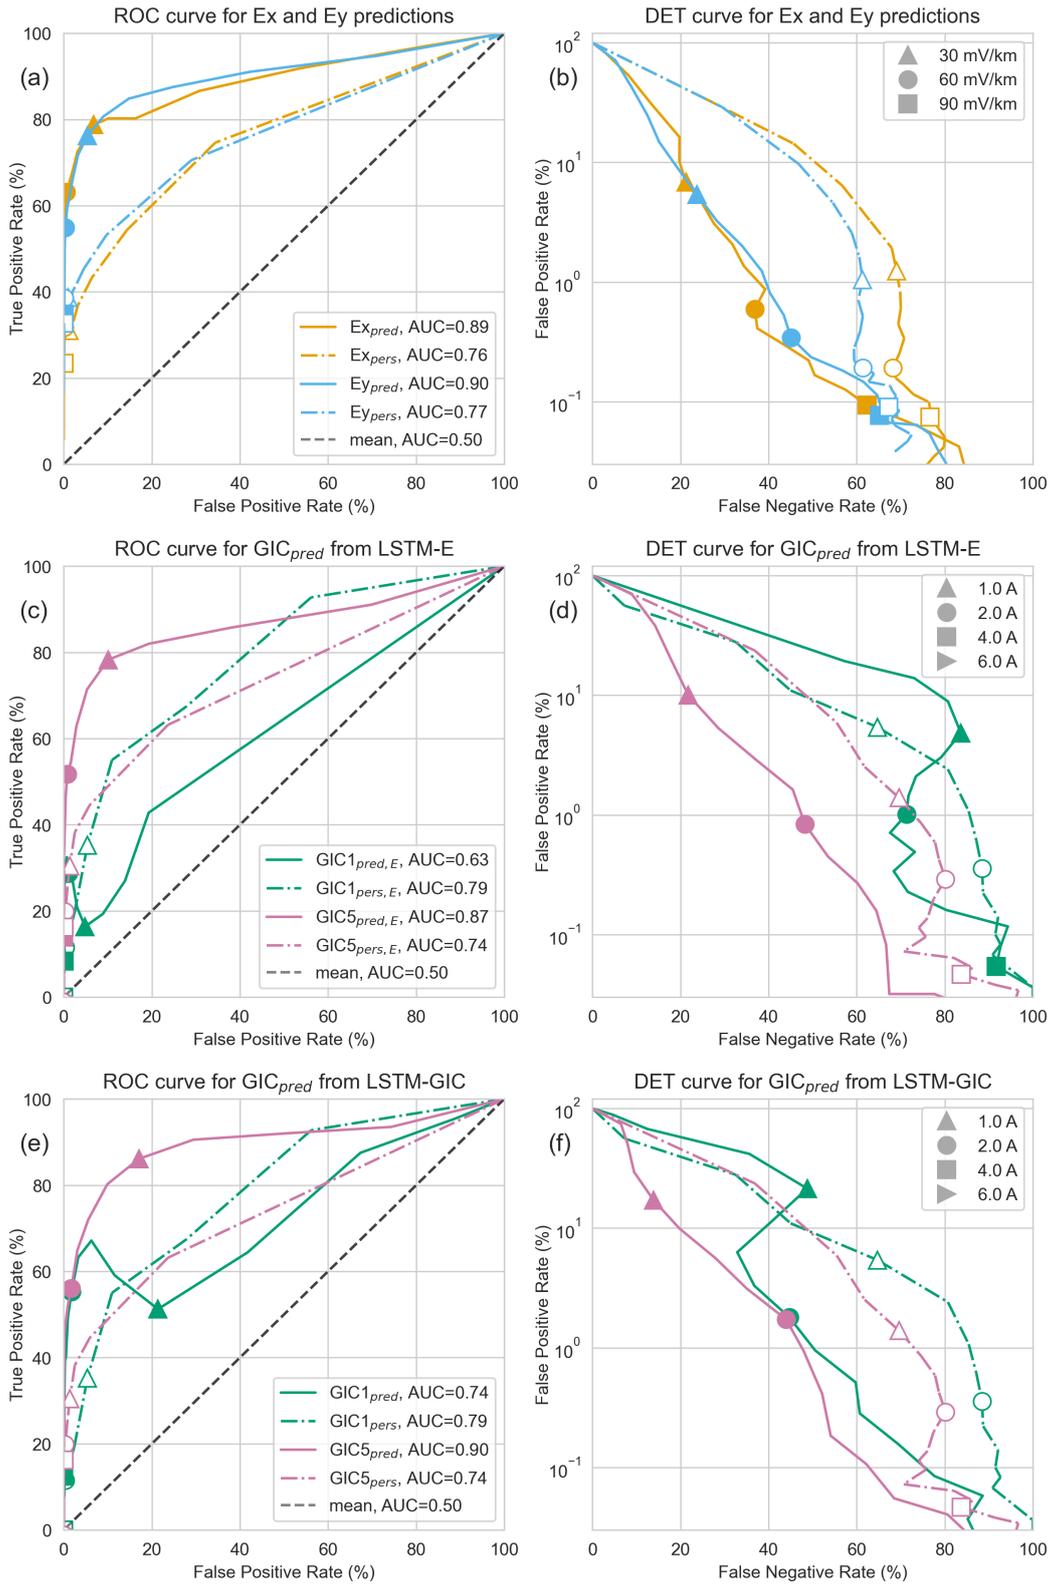


Figure 5. Receiver-operator characteristic (ROC) and detection-error tradeoff (DET) curves for three approaches: (a-b) the geoelectric field, showing the output from the LSTM-E models vs the modelled geoelectric field, (c-d) the GICs calculated from the geoelectric field predicted by LSTM-E compared to measured GICs, and (e-f) the GICs predicted by the LSTM-GIC models compared to measured GICs. SS1 and SS5 are two separate substations in the power grid from which we have measurements. The values for specific event thresholds are labelled with shapes as defined in each legend.

473 served behaviour shows a curve that keeps close to the upper left corner. The area-under-
 474 the-curve (AUC in the legend) shows good model skill as it approaches 1. On the other
 475 hand, the DET curve shows the relationship between the false negative rate (fraction of
 476 all predicted non-events that were measured events misclassified as non-events, or $FN/(FN+TN)$)
 477 and false positive rate, the number of which usually goes up as the other goes down de-
 478 pending on where the threshold for an event is set. Here, the best model behaviour is
 479 seen as the curves approach the lower left corner. It is useful in error minimisation to
 480 deduce the rate at which the FNR improves with regards to an increase in FPR rate (and
 481 vice-versa).

Table 2. Metrics from an event-based analysis of the LSTM-E models applied to the years 2000, 2001, 2017, 2019 and 2020 in a retrospective real-time mode with the model being run at 15-minute intervals. A persistence model (PERS) is included for comparison. The first four columns provide the values for the confusion matrix (where TP, FP, TN and FN are the true positives (hits), false positives, true negatives (misses) and false negatives), the probability of detection (POD), probability of false detection (POFD), Heidke Skill Score (HSS), True Skill Score (TSS), and bias (BS). The variable TH in brackets gives the event threshold used to define events and compute the metrics.

LSTM-E Model	$N_{events,obs}$	TP	FP	FN	TN	POD	POFD	HSS	TSS	BS
$E_{x,pred}(TH=30)$	3092	2436	11749	656	160506	78.8	6.8	0.26	0.72	4.6
$E_{x,pred}(TH=60)$	494	312	1038	182	173815	63.2	0.6	0.34	0.63	2.7
$E_{x,pred}(TH=90)$	175	66	164	109	175008	37.7	0.1	0.33	0.38	1.3
$E_{y,pred}(TH=30)$	2989	2279	9328	710	163030	76.2	5.4	0.29	0.71	3.9
$E_{y,pred}(TH=60)$	559	307	600	252	174188	54.9	0.3	0.42	0.55	1.6
$E_{y,pred}(TH=90)$	241	84	135	157	174971	34.9	0.1	0.36	0.35	0.9
PERS Model	$N_{events,obs}$	TP	FP	FN	TN	POD	POFD	HSS	TSS	BS
$E_{x,pers}(TH=30)$	3092	958	2128	2134	170127	31.0	1.2	0.30	0.30	1.0
$E_{x,pers}(TH=60)$	494	157	335	337	174518	31.8	0.2	0.32	0.32	1.0
$E_{x,pers}(TH=90)$	175	41	130	134	175042	23.4	0.1	0.24	0.23	1.0
$E_{y,pers}(TH=30)$	2989	1156	1804	1833	170554	38.7	1.0	0.38	0.38	1.0
$E_{y,pers}(TH=60)$	559	216	335	343	174453	38.6	0.2	0.39	0.38	1.0
$E_{y,pers}(TH=90)$	241	79	158	162	174948	32.8	0.1	0.33	0.33	1.0

482 4.1 Forecasting E_x and E_y

483 We first evaluate the LSTMs trained on the geoelectric field in terms of the root-
 484 mean-square-error (RMSE) and the Pearson's correlation coefficient (PCC). Compar-
 485 ing the LSTM-E outputs to modelled E , the RMSE values are 126 mV/km and 111 mV/km
 486 for the absolute value of E_x and E_y , while the PCC values are 0.60 and 0.61. Once the
 487 sign of E has been included, the RMSE rises to 261 mV/km and 287 mV/km, while PCC
 488 drops to 0.48 and 0.32, so we see that the model's inability to forecast the field direc-
 489 tion reliably decreases the accuracy when also considering the field direction.

490 **Table 2** presents an event-based analysis of the LSTM-E results. Multiple thresh-
 491 olds (TH) defining events were considered, and these are listed by the variable "TH" in
 492 each line (at 30, 60, and 90 mV/km, representing minor, moderate and strong geomag-
 493 netic activity). We see that the skill decreases as the threshold increases (decreasing prob-
 494 ability of detection POD and TSS), and that the LSTMs tend towards over-predicting
 495 (BS > 1). (The bias for the PERS models is always ~ 1 because the time series being

496 compared are only shifted in time and therefore almost statistically equivalent.) There
 497 are always a large number of false positives, although this remains a small fraction of
 498 the number of total data points. The LSTM-E models generally outperform the PERS
 499 approach, although the Heidke Skill Scores are occasionally smaller in the LSTMs, which
 500 implies a worse balance between false positives and true positives. As in the point-to-
 501 point values, the E_X component tends to be predicted better than the E_Y component.
 502 By evaluating the ROC and DET curves in Fig. 5 (a-b), we see that the LSTM-E mod-
 503 els outperforms persistence at all thresholds.

504 We also conducted a comparison with the results from Honkonen et al. (2018) and
 505 Lotz and Danskin (2017), where possible. While the time development of the geoelec-
 506 tric field appears better in the modelling approach in Honkonen et al. (2018), the mag-
 507 nitudes are not matched as well. An event-based analysis could not be carried out in their
 508 case due to the short time series and lack of larger events, but the RMSE and PCC val-
 509 ues for E_X and E_Y (reduced to a 15-min sampling rate) come out as 10.5 mV/km and
 510 97.8 mV/km and 0.62 and 0.25, respectively, which is better in the case of E_X but worse
 511 in the case of E_Y . Comparing to Lotz and Danskin (2017), we see similar correlations
 512 for the geoelectric field components. They found a slightly higher correlation (averaged
 513 over three stations and two storms, 0.71 for E_X and 0.53 for E_Y), although they predicted
 514 the maximum value for a longer time span (90 mins), making their approach closer to
 515 a nowcast than a forecast. The higher RMSE values seen in our study in part derive from
 516 the slightly higher levels of daily variation that is forecast even when the field is extremely
 517 quiet. Again, in both studies used as comparison we see the northward component of the
 518 geoelectric field was predicted better than the eastward component.

519 4.2 Forecasting GICs

520 The same results are presented for GICs as for the geoelectric field components in
 521 the last section. In the event-based analysis, the thresholds were set at 2, 4 and 6 A, which
 522 are roughly equivalent to the thresholds used for the electric field. **Table 3** shows the
 523 results of this analysis applied to the test data set years 2017, 2019 and 2020, while **Fig. 5**
 524 depicts the ROC and DET curves for the model output versus measured GICs. A com-
 525 parison between the LSTM-GIC output and the modelled GICs the model was trained
 526 on shows similar levels of accuracy as in LSTM-E to the geoelectric field.

527 We first look at the results for GICs calculated from the geoelectric field compo-
 528 nents predicted using the LSTM-E models. Note that while the last section mainly looked
 529 at the absolute value of the geoelectric fields, in the calculation of GICs the direction of
 530 the geoelectric field is also included, making this an additional error factor if the sign
 531 is not predicted accurately. Once the GICs have been calculated using the results from
 532 the LSTM-E models and Eq. 1, the absolute value is taken for the rest of the analysis.

533 As can be seen in **Table 3**, the GICs derived from the LSTM-E models see a con-
 534 siderable drop in accuracy in comparison to the results for E alone in Table 2. Although
 535 there were quite reasonable values for POD predicting E, the POD for GICs at the mid-
 536 range threshold (60 mV/km or 4 A) drops from around 50% in both components of E
 537 to 8% and 16% in substation SS1 and SS5. Evaluating the skill of the model for GICs
 538 at high levels is difficult because there are so few events exceeding even a minimal value
 539 of 6 A. None of these events (2 at SS1, 12 at SS5 over the three years of data) were pre-
 540 dicted using any approach.

541 In comparing the GIC predictions from the two methods (LSTM-E and LSTM-GIC),
 542 we see that the LSTM-GIC seems to perform better but the results are station-specific.
 543 The LSTM-GIC performs much better than the LSTM-E at SS1 (e.g. a POD of 55% rather
 544 than 29% and higher HSS and TSS values at a threshold of 2 A) and at a similar level
 545 at SS5. This is also reflected in a model evaluation using point-to-point metrics. The RMSE
 546 values for SS1 and SS5 predicted using LSTM-E are 0.49 A and 0.59 A, while the PCC

Table 3. Metrics from an event-based analysis of different model applied to the years 2017, 2019 and 2020 in a retrospective real-time mode with the model being run at 15-minute intervals. $GIC1_{pred,E}$ is the result from the models trained to predict the geoelectric field (LSTM-E), while $GIC1_{pred}$ is the result from the LSTM-GIC. PERS is a persistence model assuming the target (GIC) repeats itself. The first four columns provide the values for the confusion matrix (where TP, FP, TN and FN are the true positives (hits), false positives, true negatives (misses) and false negatives), the probability of detection (POD), probability of false detection (POFD), Heidke Skill Score (HSS), True Skill Score (TSS), and bias (BS). The variable TH in brackets is the event threshold used to define events and compute the metrics. "undef." refers to the HSS and TSS at TP=0, which are undefined.

LSTM-E Model	$N_{events,obs}$	TP	FP	FN	TN	POD	POFD	HSS	TSS	BS
$GIC1_{pred,E}(TH=2)$	432	124	1060	308	103697	28.7	1.0	0.15	0.28	2.7
$GIC1_{pred,E}(TH=4)$	24	2	57	22	105108	8.3	0.1	0.05	0.08	2.5
$GIC5_{pred,E}(TH=2)$	307	159	681	148	80649	51.8	0.8	0.27	0.51	2.7
$GIC5_{pred,E}(TH=4)$	43	6	13	37	81581	14.0	0.0	0.19	0.14	0.4
LSTM-GIC Model	$N_{events,obs}$	TP	FP	FN	TN	POD	POFD	HSS	TSS	BS
$GIC1_{pred}(TH=2)$	432	239	1886	193	102871	55.3	1.8	0.18	0.54	4.9
$GIC1_{pred}(TH=4)$	24	3	26	21	105139	12.5	0.0	0.11	0.12	1.2
$GIC5_{pred}(TH=2)$	307	172	1403	135	79927	56.0	1.7	0.18	0.54	5.1
$GIC5_{pred}(TH=4)$	43	7	16	36	81578	16.3	0.0	0.21	0.16	0.5
PERS Model	$N_{events,obs}$	TP	FP	FN	TN	POD	POFD	HSS	TSS	BS
$GIC1_{pers}(TH=2)$	432	50	375	382	104382	11.6	0.4	0.11	0.11	1.0
$GIC1_{pers}(TH=4)$	24	0	26	24	105139	0.0	0.0	undef.	undef.	1.1
$GIC5_{pers}(TH=2)$	307	61	237	246	81093	19.9	0.3	0.20	0.20	1.0
$GIC5_{pers}(TH=4)$	43	7	38	36	81556	16.3	0.0	0.16	0.16	1.0

547 is 0.35 and 0.67. For GICs predicted using LSTM-GIC, the RMSE values are 0.67 A and
548 0.78 A (i.e. slightly worse than LSTM-E), but the PCC is 0.56 and 0.64. The accuracy
549 between the two approaches is roughly equivalent for SS5, but using LSTM-GIC rather
550 than LSTM-E is a definite improvement for SS1 observations. Some of the reason for this
551 can be seen in **Fig. 6**. In SS1, the jumps in values computed from LSTM-E result from
552 changes in the sign of the geoelectric field components, which then cancel each other out
553 and lead to a GIC of zero. (Conversely, ignoring the sign from LSTM-E and taking the
554 absolute values to calculate the GICs in SS1 results in higher correlation and POD but
555 a far larger number of false positives, leaving this as another possibility.) In the best cases,
556 the GIC forecasts only reach a POD of 16% for GICs above a threshold of 4 A, highlight-
557 ing the difficulty in correctly predicting larger values.

558 In the ROC and DET curves in **Figure 5** panels (c-d) for GICs from LSTM-E and
559 (e-f) from LSTM-GIC, we also see some of the weak forecasting ability for SS1 primar-
560 ily represents the LSTM behaviour at low values (GICs < 1 A). At SS1, there is a mostly
561 continuous level of noise around 1 A, and the model does not predict the noise while the
562 persistence model captures it clearly. This is an example of the weakness of ROC curves,
563 where in this case only the lower left corner (showing values greater than 1 A) is of in-
564 terest to us.

565 **Figure 6** shows the forecast that would have been produced by the model (solid
566 and black dashed lines) against measurements (coloured lines) during the September 2017

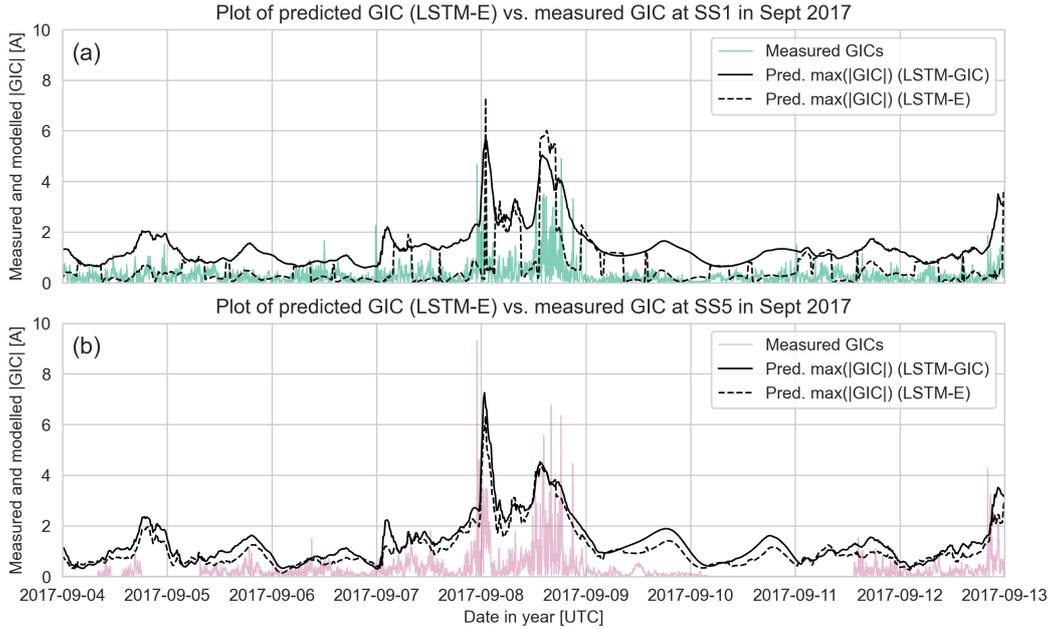


Figure 6. The LSTM-E (dashed line) and LSTM-GIC (solid line) applied to forecasts in an experimental real-time mode and compared to measurements of GICs (coloured lines) at two stations in Austria. The upper panel (a) shows results for SS1 near Vienna, while the lower panel (b) shows results for SS5 near Salzburg (with some data gaps). Although not plotted here, the maximum GIC value computed from the measurements is at the same cadence of 15 minutes to compare to the model forecasts.

567 storm. The models, particularly the LSTM-GIC approach, do a reasonable job at pre-
 568 dicting magnitudes, although the LSTM-E struggles to predict the direction, which is
 569 also important for accurate GIC prediction. The storm and the active periods are clearly
 570 captured by the forecast, and daily variations from the S_q current are forecasted oth-
 571 erwise. Note that the delayed rise in the forecast of the first peak of the storm does not
 572 indicate a timing error. A cross-correlation of the model output shows at maximum an
 573 offset in time of 10 minutes and the delay in the figure is simply a feature unique to this
 574 storm. While the exact time development of the storm is not captured well, the general
 575 scales of GICs are matched well, as is the differentiation between quiet and active times.

576 In summary, prediction of geoelectric field magnitudes can be achieved with rea-
 577 sonable accuracy (POD of at least 35% even at the highest event threshold), but the pre-
 578 diction of elevated levels of GICs proves difficult with any approach used. The LSTMs
 579 usually outperform the persistence models, except in the bias, where the persistence model
 580 has the benefit of being statistically equivalent to the data it is being compared to. The
 581 persistence model also generally has a lower POFD and higher HSS value at low thresh-
 582 olds (e.g. $TH=30$ V/km for LSTM-E) because quiet periods tend to persist over time.
 583 The LSTMs, however, outperform persistence at the higher thresholds, which are more
 584 important for forecasting purposes.

585 **5 Discussion**

586 We have attempted to forecast GICs from solar wind data using LSTMs with two
 587 different approaches. We now look at some of the reasons behind the particular difficulty
 588 in forecasting GICs.

589 Some of the low skill seen when comparing predictions to GIC measurements is down
 590 to four reasons, mostly related to our GIC data: firstly, there is noise in the GIC obser-
 591 vations, particularly at SS1, which has a consistent level of 1 A noise during the day -
 592 this is not predicted by the model. Secondly, GIC observations until 2021 had a max-
 593 imum cutoff point of 3.4 A in the positive direction, removing some peaks from our event
 594 list, and these have not been accounted for. Thirdly, the model struggles to predict the
 595 direction of the geoelectric field values, which are likely driven by smaller-scale ionospheric
 596 currents (Dimmock et al., 2020). Fourthly, as noted in Sec. 2.3, the peaks of observed
 597 GICs are often underestimated by geophysical modelling, meaning peaks in the GIC mea-
 598 surements after the cut-off level was removed were often much larger than modelled. This
 599 is a problem related to the geoelectric field modelling that may affect the LSTM’s abil-
 600 ity to learn the problem due to insufficient accuracy in the field modelling. While minute
 601 cadence data does capture most of the variability in the GICs, the lack of higher frequency
 602 content appears to be the primary cause of underestimated peaks, a problem discussed be-
 603 fore in Grawe et al. (2018) and recently for the specific problem of GIC estimates in Beggan
 604 et al. (2021). As such, it is not surprising that the LSTMs tend to underestimate the ac-
 605 tual GICs, and a correction would have to be applied to the target data to account for
 606 this.

607 Outside of the data-specific problems, there are also some timing errors, meaning
 608 some peaks arrived slightly later or earlier than they were observed, and as such are not
 609 logged as correct predictions even though an event threshold was crossed.

610 In an application of the model in operations, one caveat is that the maximum possi-
 611 ble forecast is 200 mV/km due to a self-imposed limit to improve the model’s ability
 612 to learn. We assume that in practise, this would be negligible because all values above
 613 a certain level (e.g. 100 mV/km) would be of interest, regardless of how large they be-
 614 come. As also discussed in Wintoft et al. (2016), the scale of geomagnetic variations dur-
 615 ing extreme events can theoretically become so large that it is effectively unbounded for
 616 the purpose of this discussion. In the future, this 200 mV/km limit could be improved
 617 on by training a model specifically for large value forecasting, which can be switched to
 618 if the original model forecasts $E > 150$ mV/km.

619 In an ideal case, a forecasting model would be developed while taking a cost-loss
 620 analysis (Murphy, 1977) such as that used in a space weather context in Owens et al.
 621 (2014) into consideration. In the case of network protection, this is a very complex sce-
 622 nario due to the varying impacts and costs associated with transformer damage or power
 623 grid outage, many of which are currently nearly impossible to estimate. This is some-
 624 thing that can hopefully be developed further as studies into GIC risk progress (Eastwood
 625 et al., 2018).

626 Another, more general problem in forecasting any measure of ground geomagnetic
 627 activity from solar wind measurements without further input from the magnetosphere-
 628 ionosphere system is that not all geomagnetic variations are driven by the solar wind di-
 629 rectly (see e.g. Kamide et al., 1998; Eastwood et al., 2015). Many of the ground vari-
 630 ations, particularly at shorter timescales (Alberti et al., 2017), are not directly driven
 631 by the solar wind but are instead the consequence of other processes being triggered. These
 632 can include complex magnetospheric dynamics such as reconnection in the magnetotail,
 633 as well as random, chaotic processes. Such processes can not be related in detail through
 634 our model, which is essentially a coupling function from the solar wind at the bow shock
 635 to the geoelectric field in Austria. Some of the dynamics will be represented to some de-

636 gree, but it is difficult to ascertain exactly which in a black-box machine learning model.
 637 A further difficulty in improving predictions lies in the fact that GICs can only be cal-
 638 culated accurately with knowledge of magnetic field variations at timescales of seconds
 639 (Grawe et al., 2018), ideally, and the LSTM must make approximations of what kind of
 640 variations are expected due to the conditions rather than deriving the variations precisely.
 641 Although the machine learning approach described here works at a basic level and could
 642 be more promising than forecasts of dB/dt alone, to create a model that can also account
 643 for complex magnetospheric processes it would need to be coupled with either data from
 644 space-borne monitors observing the Earth’s magnetosphere, more complex physical mod-
 645 els of magnetospheric behaviour (developing a so-called grey-box model as recommended
 646 in Camporeale (2019), for example), or both.

647 The calculations and measurements of the GICs shown in this study are for a spec-
 648 ific grid configuration, even though the power grid is continually being upgraded and
 649 changed. These changes can have large effects on individual GIC scales over long time
 650 ranges. The results shown in Table 1 extend far into the past, for which we do not have
 651 a detailed history of grid changes, so the values listed could have been much smaller or
 652 much larger depending on how the grid was set up. For the LSTM predictions, we have
 653 conducted our analysis with the comparison to measurements over a considerably shorter
 654 time range of a few years, where the grid has not changed to any great degree, but the
 655 predictions may not be valid in the future for a different grid configuration. In this case,
 656 a new fit would need to be found for Eq. 1, and either the LSTM-GIC model would need
 657 to be retrained on the updated GIC data, or the GIC values could be calculated anew
 658 from the otherwise unchanged LSTM-E output.

659 Our aim was to develop a model that can provide useful forecasts for power grid
 660 operators by providing estimates of the scales of GICs. The difference between this and
 661 former studies such as Lotz et al. (2017) and Honkonen et al. (2018), who also predicted
 662 ground geoelectric fields from solar wind data, is that we have approached the problem
 663 with a new tool (a recurrent neural network) and have been able to forecast GICs di-
 664 rectly along with the geoelectric field, with the results compared to measured GICs. We
 665 have had some success, particularly with forecasting the geoelectric field, and have tried
 666 forecasting substation-specific GICs for the first time, but there are still many problems
 667 to be addressed to turn this method into a useful forecast.

668 6 Summary

669 We have developed a machine learning approach to forecast GICs in Austria. Us-
 670 ing data from the past 26 years and the 2003 Halloween storm as a case study, we ar-
 671 gued that forecasts of dB/dt alone, which have been the focus of most past studies, are
 672 not sufficient to make actionable GIC forecasts.

673 From this initial analysis, we set out to forecast maximum expected GICs (over a
 674 forty minute window) either directly for specific substations in the power grid or more
 675 generally from forecasts of the regional geoelectric field components. From a small set
 676 of initial machine learning approaches, an LSTM (recurrent neural network) with an At-
 677 tention mechanism showed the most promise in forecasting skill and this was developed
 678 into a more complex approach.

679 A selection of models were trained on 21 years of geoelectric field values modelled
 680 from geomagnetic variations at the geomagnetic observatory in Fürstenfeldbruck close
 681 to Austria. In the first method, two recurrent neural networks or LSTMs were trained
 682 to predict the northward and eastward modelled geoelectric field components and com-
 683 pute the specific substation GICs using a linear equation. In the second method, an LSTM
 684 was trained to predict modelled GICs at two substations, which we know correlate very
 685 well with the measurements. Five years of data were reserved for testing and evaluat-

686 ing the model. The results were compared to DC measurements at two substations in
 687 the Austrian power grid.

688 The LSTM model worked with reasonable success when predicting the geoelectric
 689 field modelled from geomagnetic variations, although translating this success into good
 690 GIC forecasts proved difficult. It was possible, however, to outperform a model that sim-
 691 ply takes the last observed GICs to forecast future values.

692 We conclude that forecasting the GICs observed in the power grid from solar wind
 693 data measured at L1 is a difficult task, even when the forecasting model does a reason-
 694 able job of forecasting the geoelectric field components or modelled GIC. There are many
 695 ways to improve the modelling in the future, including using higher-resolution magnetic
 696 field measurements (or applying a correction to the modelled geoelectric field before train-
 697 ing) to more accurately estimate the peak geoelectric field and GIC values, and by in-
 698 cluding information on the development of the magnetosphere during storm times.

699 Although this study has looked specifically at a mid-latitude region, where geomag-
 700 netic variations and GICs are not as large as those seen in higher latitude regions such
 701 as Scandinavia, we have been able to compare model output directly to measurements
 702 and expect that the conclusions drawn will also be valid for other regions with GICs at
 703 different scales.

704 A lower-resolution version of the LSTM-E model will be coupled with the PRED-
 705 STORM solar wind forecast (Bailey et al., 2020), which provides forecasts of the ambi-
 706 ent solar wind a few days in advance, based on either a recurrence model or data from
 707 a spacecraft east of the Sun-Earth line such as STEREO or a future mission to the La-
 708 grange 5 point. We also plan in the future to integrate methods on solar wind B_z fore-
 709 casting (Reiss et al., 2021) or CME flux rope modelling (Weiss et al., 2021) to advance
 710 our capabilities in GIC forecasting for any type of solar wind structures.

711 7 Data Availability

- 712 • INTERMAGNET data for FUR and WIC:
 713 <https://intermagnet.org/data-donnee/download-eng.php>
- 714 • OMNI data: https://spdf.gsfc.nasa.gov/pub/data/omni/high_res_omni/
- 715 • Open source code for this work (in Python 3 and Jupyter Notebook form):
 716 <https://doi.org/10.5281/zenodo.5704715>
- 717 • Exact details on the LSTM structure and hyperparameters used for training can
 718 be found in the supporting information for this study.
- 719 • A subset of the data set used to derive the results, namely the the GIC observa-
 720 tions and model forecasts used to produce Figure 6, have also been included in the
 721 supporting information and saved in an online repository:
 722 <https://doi.org/10.6084/m9.figshare.19102772.v1>

723 Acknowledgments

724 The results presented in this paper rely on the data collected at the Conrad Observa-
 725 tory (ZAMG) in Austria and at Fürstfeldbruck (LMU), Germany. We thank the Ludwig-
 726 Maximilians-Universität München for supporting its operation and INTERMAGNET
 727 for promoting high standards of magnetic observatory practice (www.intermagnet.org).
 728 The data used in this study is publicly available (with the exception of the measurements
 729 of GICs in Austria), and details on where to find the data can be found in Section 7. An
 730 excerpt of GIC measurements and model output has been provided in the Supporting
 731 Information, and we thank both the TU Graz and Austrian Power Grid for supporting
 732 these measurements. We thank Ilja Honkonen for providing the results from his 2018 study
 733 for comparison. R.L.B., C.M., M.A.R., and A.J.W. thank the Austrian Science Fund (FWF)

734 for research funding from projects P31659-N27 and P31521-N27. C.D.B. was funded un-
 735 der UK Natural Environment Research Council Grant NE/P017231/1 “Space Weather
 736 Impact on Ground-based Systems (SWIGS)”. We thank the two anonymous reviewers
 737 for their careful readings of this manuscript and helpful suggestions.

738 References

- 739 m, A., Pracser, E., & Wesztergom, V. (2012). Estimation of the electric resistivity
 740 distribution (EURHOM) in the European lithosphere in the frame of the
 741 EURISGIC WP2 project. *Acta Geodaetica et Geophysica Hungarica*, *47*(4),
 742 377-387.
- 743 Albert, D., Schachinger, P., Bailey, R. L., Renner, H., & Achleitner, G. (2022). Anal-
 744 ysis of long-term gic measurements in transformers in austria. *Space Weather*,
 745 *20*(1), e2021SW002912.
- 746 Alberti, T., Consolini, G., Lepreti, F., Laurenza, M., Vecchio, A., & Carbone, V.
 747 (2017). Timescale separation in the solar wind-magnetosphere coupling during
 748 st. patrick’s day storms in 2013 and 2015. *Journal of Geophysical Research:*
 749 *Space Physics*, *122*(4), 4266-4283.
- 750 Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly
 751 learning to align and translate. In *3rd international conference on learning rep-*
 752 *resentations*.
- 753 Bailey, R. L., Halbedl, T. S., Schattauer, I., Achleitner, G., & Leonhardt, R. (2018).
 754 Validating GIC models with measurements in Austria: evaluation of accuracy
 755 and sensitivity to input parameters. *Space Weather*.
- 756 Bailey, R. L., Halbedl, T. S., Schattauer, I., Romer, A., Achleitner, G., Beggan,
 757 C. D., ... Leonhardt, R. (2017). Modelling geomagnetically induced currents
 758 in midlatitude Central Europe using a thin-sheet approach. *Annales*
 759 *Geophysicae*, *35*(3), 751.
- 760 Bailey, R. L., Leonhardt, R., Mostl, C., Beggan, C., Reiss, M. A., Bhaskar, A., &
 761 Weiss, A. J. (n.d.). Forecasting gics and geoelectric fields from solar wind data
 762 using lstms: application in austria. *Space Weather*, *n/a*(n/a), e2021SW002907.
 763 Retrieved from [https://agupubs.onlinelibrary.wiley.com/doi/](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2021SW002907)
 764 [abs/10.1029/2021SW002907](https://doi.org/10.1029/2021SW002907) (e2021SW002907 2021SW002907) doi:
 765 <https://doi.org/10.1029/2021SW002907>
- 766 Bailey, R. L., Mostl, C., Reiss, M. A., Weiss, A. J., Amerstorfer, U. V., Amerstorfer,
 767 T., ... Leonhardt, R. (2020). Prediction of Dst during solar minimum using in
 768 situ measurements at 15. *Space Weather*, *18*(5), e2019SW002424.
- 769 Baker, D. N., Li, X., Pulkkinen, A., Ngwira, C. M., Mays, M. L., Galvin, A. B., &
 770 Simunac, K. D. C. (2013, October). A major solar eruptive event in July 2012:
 771 Defining extreme space weather scenarios. *Space Weather*, *11*(10), 585-591.
 772 doi: 10.1002/swe.20097
- 773 Barbosa, C., Alves, L., Caraballo, R., Hartmann, G. A., Papa, A. R., & Pirjola,
 774 R. J. (2015). Analysis of geomagnetically induced currents at a low-latitude
 775 region over the solar cycles 23 and 24: comparison between measurements and
 776 calculations. *Journal of Space Weather and Space Climate*, *5*, A35.
- 777 Beggan, C. D., Richardson, G. S., Baillie, O., Hubert, J., & Thomson, A. W. (2021).
 778 Geoelectric field measurement, modelling and validation during geomagnetic
 779 storms in the uk. *Journal of Space Weather and Space Climate*, *11*, 37.
- 780 Bhaskar, A., & Vichare, G. (2019). Forecasting of SYMH and ASYH indices for
 781 geomagnetic storms of solar cycle 24 including St. Patrick’s day, 2015 storm
 782 using NARX neural network. *J. Space Weather Space Clim.*, *9*, A12.
- 783 Bloomfield, D. S., Higgins, P. A., McAteer, R. T. J., & Gallagher, P. T. (2012,
 784 Mar). Toward Reliable Benchmarking of Solar Flare Forecasting Methods. *The*
 785 *Astrophysical Journal Letters*, *747*(2), L41.

- 786 Bolduc, L. (2002). GIC observations and studies in the Hydro-Québec power system.
787 *Journal of Atmospheric and Solar-Terrestrial Physics*, *64*(16), 1793–1802.
- 788 Boteler, D. H., & Pirjola, R. J. (2017). Modeling geomagnetically induced currents.
789 *Space Weather*, *15*(1), 258–276.
- 790 Boteler, D. H., Pirjola, R. J., & Nevanlinna, H. (1998). The effects of geomagnetic
791 disturbances on electrical systems at the Earth’s surface. *Advances in Space*
792 *Research*, *22*, 17–27.
- 793 Butala, M. D., Kazerooni, M., Makela, J. J., Kamalabadi, F., Gannon, J. L., Zhu,
794 H., & Overbye, T. J. (2017). Modeling geomagnetically induced currents from
795 magnetometer measurements: Spatial scale assessed with reference measure-
796 ments. *Space Weather*, *15*(10), 1357–1372.
- 797 Camporeale, E. (2019). The challenge of machine learning in space weather: Now-
798 casting and forecasting. *Space Weather*, *17*(8), 1166–1207.
- 799 Caraballo, R., González-Esparza, J. A., Sergeeva, M., & Pacheco, C. R. (2020).
800 First GIC estimates for the Mexican power grid. *Space Weather*, *18*(2),
801 e2019SW002260.
- 802 Chave, A. D., & Jones, A. G. (2012). *The magnetotelluric method: Theory and prac-*
803 *tice*. Cambridge University Press.
- 804 Dimmock, A. P., Rosenqvist, L., Welling, D. T., Viljanen, A., Honkonen, I., Boynton,
805 R. J., & Yordanova, E. (2020). On the regional variability of db/dt and
806 its significance to gic. *Space Weather*, *18*(8), e2020SW002497.
- 807 Eastwood, J., Hapgood, M., Biffis, E., Benedetti, D., Bisi, M., Green, L., . . . Bur-
808 nett, C. (2018). Quantifying the economic value of space weather forecasting
809 for power grids: An exploratory study. *Space Weather*, *16*(12), 2052–2067.
- 810 Eastwood, J., Hietala, H., Toth, G., Phan, T., & Fujimoto, M. (2015). What con-
811 trols the structure and dynamics of Earth’s magnetosphere? *Space Science Re-*
812 *views*, *188*(1), 251–286.
- 813 Galassi, A., Lippi, M., & Torroni, P. (2020). Attention in natural language process-
814 ing. *IEEE Transactions on Neural Networks and Learning Systems*.
- 815 Gaunt, C. T., & Coetzee, G. (2007). Transformer failures in regions incorrectly con-
816 sidered to have low GIC-risk. In *Power tech, 2007 IEEE Lausanne* (p. 807–812).
817 IEEE.
- 818 Gil, A., Modzelewska, R., Moskwa, S., Siluszyk, A., Siluszyk, M., Wawrzynczak,
819 A., & Zakrzewska, S. (2019). Does time series analysis confirms the relation-
820 ship between space weather effects and the failures of electrical grids in south
821 poland? *Journal of Mathematics in Industry*, *9*(1), 1–16.
- 822 Gopalswamy, N., Yashiro, S., Liu, Y., Michalek, G., Vourlidas, A., Kaiser, M. L.,
823 & Howard, R. A. (2005, September). Coronal mass ejections and other
824 extreme characteristics of the 2003 October–November solar eruptions.
825 *Journal of Geophysical Research (Space Physics)*, *110*(A9), A09S15. doi:
826 10.1029/2004JA010958
- 827 Grawe, M. A., Makela, J. J., Butala, M. D., & Kamalabadi, F. (2018). The impact
828 of magnetic field temporal sampling on modeled surface electric fields. *Space*
829 *Weather*, *16*(11), 1721–1739.
- 830 Heidke, P. (1926). Berechnung des erfolges und der güte der windstärkevorhersagen
831 im sturmwarnungsdienst. *Geografiska Annaler*, *8*(4), 301–349.
- 832 Honkonen, I., Kuvshinov, A., Rastätter, L., & Pulkkinen, A. (2018). Predicting
833 global ground geoelectric field with coupled geospace and three-dimensional
834 geomagnetic induction models. *Space Weather*, *16*(8), 1028–1041.
- 835 Juusola, L., Viljanen, A., van de Kamp, M., Tanskanen, E. I., Vanhamäki, H., Par-
836 tamies, N., & Kauristie, K. (2015). High-latitude ionospheric equivalent
837 currents during strong space storms: Regional perspective. *Space Weather*,
838 *13*(1), 49–60.
- 839 Kamide, Y., Baumjohann, W., Daglis, I. A., Gonzalez, W. D., Grande, M., Jose-
840 lyn, J. A., . . . Vasyliunas, V. M. (1998). Current understanding of magnetic

- 841 storms: Storm-substorm relationships. *Journal of Geophysical Research: Space*
 842 *Physics*, 103(A8), 17705-17728.
- 843 Keesee, A. M., Pinto, V., Coughlan, M., Lennox, C., Mahmud, M. S., & Connor,
 844 H. K. (2020). Comparison of deep learning techniques to model connections
 845 between solar wind and ground magnetic perturbations. *Frontiers in Astron-*
 846 *omy and Space Sciences*, 7, 72.
- 847 Kelbert, A. (2020). The role of global/regional earth conductivity models in natural
 848 geomagnetic hazard mitigation. *Surv Geophys*, 41, 115–166.
- 849 Liu, Y. D., Luhmann, J. G., Kajdič, P., Kilpua, E. K. J., Lugaz, N., Nitta, N. V., ...
 850 Galvin, A. B. (2014, March). Observations of an extreme storm in interplan-
 851 etary space caused by successive coronal mass ejections. *Nature Communica-*
 852 *tions*, 5, 3481. doi: 10.1038/ncomms4481
- 853 Lotz, S., & Cilliers, P. (2015). A solar wind-based model of geomagnetic field fluc-
 854 tuations at a mid-latitude station. *Advances in Space Research*, 55(1), 220 -
 855 230.
- 856 Lotz, S., & Danskin, D. W. (2017). Extreme value analysis of induced geoelectric
 857 field in South Africa. *Space Weather*, 15(10), 1347-1356.
- 858 Lotz, S., Heyns, M., & Cilliers, P. J. (2017). Regression-based forecast model of in-
 859 duced geoelectric field. *Space Weather*, 15(1), 180-191.
- 860 Lu, J., Peng, Y., Wang, M., Gu, S., & Zhao, M. (2016). Support Vector Machine
 861 combined with distance correlation learning for Dst forecasting during intense
 862 geomagnetic storms. *Planetary and Space Science*, 120, 48 - 55.
- 863 Molinski, T. S. (2002). Why utilities respect geomagnetically induced currents. *Jour-*
 864 *nal of atmospheric and solar-terrestrial physics*, 64(16), 1765–1778.
- 865 Murphy, A. H. (1977). The value of climatological, categorical and probabilistic
 866 forecasts in the cost-loss ratio situation. *Monthly Weather Review*, 105(7),
 867 803–816.
- 868 Ngwira, C. M., Pulkkinen, A., Leila Mays, M., Kuznetsova, M. M., Galvin, A.,
 869 Simunac, K., ... Glocer, A. (2013). Simulation of the 23 July 2012 extreme
 870 space weather event: What if this extremely rare CME was Earth directed?
 871 *Space Weather*, 11(12), 671–679.
- 872 Ngwira, C. M., Pulkkinen, A. A., Bernabeu, E., Eichner, J., Viljanen, A., & Crow-
 873 ley, G. (2015). Characteristics of extreme geoelectric fields and their possible
 874 causes: Localized peak enhancements. *Geophysical Research Letters*, 42(17),
 875 6916-6921.
- 876 Owens, M. J., Horbury, T. S., Wicks, R. T., McGregor, S. L., Savani, N. P., &
 877 Xiong, M. (2014). Ensemble downscaling in coupled solar wind-magnetosphere
 878 modeling for space weather forecasting. *Space Weather*, 12(6), 395-405.
- 879 Owens, M. J., Lockwood, M., Barnard, L. A., Scott, C. J., Haines, C., & Macneil,
 880 A. (2021, May). Extreme Space-Weather Events and the Solar Cycle. *Solar*
 881 *Physics*, 296(5), 82. doi: 10.1007/s11207-021-01831-3
- 882 Price, P. R. (2002). Geomagnetically induced current effects on transformers. *IEEE*
 883 *Transactions on Power Delivery*, 17(4), 1002-1008.
- 884 Pulkkinen, A., Bernabeu, E., Thomson, A., Viljanen, A., Pirjola, R., Boteler, D., ...
 885 MacAlester, M. (2017). Geomagnetically induced currents: Science, engineer-
 886 ing, and applications readiness. *Space Weather*, 15(7), 828-856.
- 887 Pulkkinen, A., Hesse, M., Habib, S., Van der Zel, L., Damsky, B., Policelli, F., ...
 888 Creamer, E. (2010). Solar shield: forecasting and mitigating space weather
 889 effects on high-voltage power transmission systems. *Natural hazards*, 53(2),
 890 333–345.
- 891 Pulkkinen, A., Hesse, M., Kuznetsova, M., & Rastätter, L. (2007). First-principles
 892 modeling of geomagnetically induced electromagnetic fields and currents from
 893 upstream solar wind to the surface of the Earth. *Annales Geophysicae*, 25(4),
 894 881–893.
- 895 Pulkkinen, A., Pirjola, R., & Viljanen, A. (2007). Determination of ground con-

- ductivity and system parameters for optimal modeling of geomagnetically induced current flow in technological systems. *Earth, planets and space*, 59(9), 999–1006.
- Pulkkinen, A., Rastatter, L., Kuznetsova, M., Singer, H., Balch, C., Weimer, D., ... Weigel, R. (2013, June). Community-wide validation of geospace model ground magnetic field perturbation predictions to support model transition to operations. *Space Weather*, 11(6), 369-385.
- Pulkkinen, A., Taktakishvili, A., Odstrcil, D., & Jacobs, W. (2009). Novel approach to geomagnetically induced current forecasts based on remote solar observations. *Space Weather*, 7(8).
- Pulkkinen, A., Viljanen, A., & Pirjola, R. (2006). Estimation of geomagnetically induced current levels from different input data. *Space Weather*, 4(8).
- Reiss, M. A., Möstl, C., Bailey, R. L., Rüdissler, H. T., Amerstorfer, U. V., Amerstorfer, T., ... Windisch, A. (2021, August). Machine learning for predicting the Bz magnetic field component from upstream in situ observations of solar coronal mass ejections. *arXiv e-prints*, arXiv:2108.04067.
- Shepherd, S. G. (2014). Altitude-adjusted corrected geomagnetic coordinates: Definition and functional approximations. *Journal of Geophysical Research: Space Physics*, 119(9), 7501-7521.
- Sun, R., & Balch, C. (2019). Comparison between 1-D and 3-D geoelectric field methods to calculate geomagnetically induced currents: A case study. *IEEE Transactions on Power Delivery*, 34(6), 2163-2172.
- Svanda, Michal, Mourenas, Didier, Zertová, Karla, & Výbost'oková, Tatiana. (2020). Immediate and delayed responses of power lines and transformers in the Czech electric power grid to geomagnetic storms. *J. Space Weather Space Clim.*, 10, 26.
- Tasistro-Hart, A., Grayver, A., & Kuvshinov, A. (2021). Probabilistic geomagnetic storm forecasting via deep learning. *Journal of Geophysical Research: Space Physics*, 126(1), e2020JA028228.
- Torta, J. M., Serrano, L., Regué, J. R., Sánchez, A. M., & Roldán, E. (2012). Geomagnetically induced currents in a power grid of northeastern Spain. *Space Weather*, 10(6).
- Viljanen, A., Pirjola, R., Prácer, E., Katkalov, J., & Wik, M. (2014). Geomagnetically induced currents in Europe - modelled occurrence in a continent-wide power grid. *J. Space Weather Space Clim.*, 4, A09.
- Weigel, R. S. (2017). A comparison of methods for estimating the geoelectric field. *Space Weather*, 15(2), 430-440.
- Weiss, A. J., Möstl, C., Amerstorfer, T., Bailey, R. L., Reiss, M. A., Hinterreiter, J., ... Bauer, M. (2021). Analysis of coronal mass ejection flux rope signatures using 3DCORE and approximate Bayesian Computation. *The Astrophysical Journal Supplement Series*, 252(1), 9.
- Welling, D. T., Ngwira, C. M., Opgenoorth, H., Haiducek, J. D., Savani, N. P., Morley, S. K., ... Liemohn, M. (2018). Recommendations for next-generation ground magnetic perturbation validation. *Space Weather*, 16(12), 1912-1920.
- Wilks, D. S. (2011). *Statistical methods in the atmospheric sciences*. Amsterdam; Boston: Elsevier Academic Press.
- Wintoft, P. (2005, July). Study of the solar wind coupling to the time difference horizontal geomagnetic field. *Annales Geophysicae*, 23(5), 1949-1957.
- Wintoft, P., Viljanen, A., & Wik, M. (2016). Extreme value analysis of the time derivative of the horizontal magnetic field and computed electric field. *Annales Geophysicae*, 34(4), 485–491.
- Wintoft, P., & Wik, M. (2021). Exploring three recurrent neural network architectures for geomagnetic predictions. *Frontiers in Astronomy and Space Sciences*, 8, 72.
- Wintoft, P., Wik, M., & Viljanen, A. (2015). Solar wind driven empirical forecast

951 models of the time derivative of the ground magnetic field. *Journal of Space*
952 *Weather and Space Climate*, 5, A7.
953 Zhang, J. J., Wang, C., Sun, T. R., Liu, C. M., & Wang, K. R. (2015). GIC due
954 to storm sudden commencement in low-latitude high-voltage power network in
955 China: Observation and simulation. *Space Weather*, 13(10), 643-655.