ADVANCING EARTH AND SPACE SCIENCES

Space Weather

RESEARCH ARTICLE
10.1029/2025SW004364

Special Collection:
The Space Weather Research to Operation to Research (R2O2R) Pipeline(s): Progress, Challenges and Prospects

Implementing an Operational Cloud-Based Now- and Forecasting System for Space Weather Ground Effects in the UK

C. D. Beggan[1], J. P. Eastwood[2], J. W. B. Eggington[2], C. Forsyth[3], M. P. Freeman[4], E. Henley[5], M. Heyns[2], J. Hübert[1], D. R. Jackson[5], A. T. LaMoury[2], G. S. Richardson[1], A. W. Smith[6], and A. W. P. Thomson[1]

[1]British Geological Survey, Edinburgh, UK, [2]Imperial College, London, UK, [3]MSSL, University College, London, UK, [4]British Antarctic Survey, Cambridge, UK, [5]Met Office, Exeter, UK, [6]University of Northumbria, Newcastle, UK

Key Points:

- Real time operational nowcast of geomagnetically induced currents in the UK
- Machine learning forecasts of magnetic field change from L1 solar wind measurements
- Implementation of real-time magnetospheric magnetohydrodynamic modeling code using L1 measurements

Correspondence to:
C. D. Beggan,
ciar@bgs.ac.uk

Citation:
Beggan, C. D., Eastwood, J. P., Eggington, J. W. B., Forsyth, C., Freeman, M. P., Henley, E., et al. (2025). Implementing an operational cloud-based now- and forecasting system for space weather ground effects in the UK. *Space Weather*, *23*, e2025SW004364. https://doi.org/10.1029/2025SW004364

Received 24 JAN 2025
Accepted 30 APR 2025

Author Contributions:
Conceptualization: C. D. Beggan, J. P. Eastwood, C. Forsyth, M. P. Freeman, E. Henley, D. R. Jackson, G. S. Richardson, A. W. P. Thomson
Data curation: J. Hübert
Formal analysis: C. D. Beggan, C. Forsyth, M. P. Freeman, M. Heyns, A. T. LaMoury, G. S. Richardson, A. W. Smith
Funding acquisition: C. D. Beggan, J. P. Eastwood, C. Forsyth, M. P. Freeman, A. W. P. Thomson
Investigation: C. D. Beggan, M. Heyns, J. Hübert, G. S. Richardson
Methodology: C. D. Beggan, J. P. Eastwood, J. W. B. Eggington,

Abstract The enhanced variation of the magnetic field during severe to extreme geomagnetic storms induces a large geoelectric field in the subsurface. Grounded infrastructure can be susceptible to geomagnetically induced currents (GICs) during these events. Modeling the effect in real-time and forecasting the magnitude of GICs are important for allowing operators of critical infrastructure to make informed decisions on potential impacts. As part of the UK-funded Space Weather Innovation, Measurement, Modeling and Risk (SWIMMR) program, we implemented nine research-level models into operational codes capable of running consistently and robustly to produce estimates of GICs in the Great Britain high voltage power transmission network, the high pressure gas pipeline network and the railway network. To improve magnetic coverage and geoelectric field modeling accuracy, three new variometer sites were installed in the UK and a 3 year campaign of magnetotelluric measurements at 53 sites was undertaken. The models rely on real-time ground observatory data and solar wind data from satellites at the L1 Lagrange point. A mixture of empirical machine learning and numerical magnetohydrodynamic models are used for forecasting. In addition to nowcast capabilities, contextual information on the likelihood of substorms, sudden commencements and large rates of change of the magnetic field were developed. The final nowcast and forecast codes were implemented in a cloud-based environment using modern software tools and practices. We describe the process to move from research to operations (R2O).

Plain Language Summary The Earth's magnetic field changes rapidly, by up to 10% in magnitude, during severe geomagnetic storms, inducing a geoelectric field in the subsurface. The geoelectric field can equalize across large distances, by flowing into low resistance infrastructure including the earthing points of high voltage transformers, long-distance gas pipelines and electrified railway lines. We have created new operational codes from scientific models to allow estimates of the effects on this critical technology. Using real-time measurements of the magnetic field from UK geomagnetic observatories in conjunction with conductivity maps of the subsurface allows us to calculate the geoelectric field. From the known electrical resistance properties of the grounded networks we can compute the instantaneous currents in them. In addition, we have developed two methods to forecast changes of the ground magnetic field using machine learning and physics-based magnetospheric models. These are driven by satellite measurements of the solar wind at the L1 Lagrange point. There are also other codes which estimate the near-future activity of the magnetic field to help forecasters gain a richer understanding of the environment. We describe how this new capability is set up on a cloud-based web service to ingest magnetic and solar wind data and compute a series of space weather outputs every 5 min.

1. Introduction

The UK Space Weather Innovation, Measurement, Modeling and Risk (SWIMMR) project was a £20 million, 4-year research-to-operations (R2O) program which ran between 2020 and 2024 with the aim of improving national capabilities in real-time monitoring and prediction of space weather hazards and with the goals of allowing actionable information to be provided to stakeholders as quickly as possible. Space weather poses significant risks to much of the technical infrastructure modern society relies on in daily life, from satellite communications to aviation to continuous electrical power, and as such is recorded in the UK's National Risk Register with an impact level just below that of a global pandemic (e.g., Cannon, 2013; Hapgood et al., 2021). The program was funded by UK Research and Innovation (UKRI) and sought to move space weather hazard capabilities from research level

C. Forsyth, M. P. Freeman, M. Heyns, J. Hübert, A. T. LaMoury, G. S. Richardson, A. W. Smith
Project administration: C. D. Beggan, J. P. Eastwood, C. Forsyth, M. P. Freeman, J. Hübert, D. R. Jackson, A. W. P. Thomson
Resources: C. D. Beggan
Software: C. D. Beggan, J. W. B. Eggington, M. P. Freeman, M. Heyns, J. Hübert, A. T. LaMoury, G. S. Richardson, A. W. Smith
Supervision: C. D. Beggan, J. P. Eastwood, C. Forsyth, M. P. Freeman, E. Henley, J. Hübert, D. R. Jackson, A. W. P. Thomson
Validation: C. D. Beggan, J. W. B. Eggington, M. Heyns, J. Hübert, A. T. LaMoury, G. S. Richardson, A. W. Smith
Visualization: M. Heyns, J. Hübert, A. T. LaMoury, G. S. Richardson, A. W. Smith
Writing – original draft: C. D. Beggan, E. Henley, M. Heyns, J. Hübert, D. R. Jackson, A. T. LaMoury, G. S. Richardson, A. W. Smith
Writing – review & editing: J. W. B. Eggington, C. Forsyth, M. P. Freeman, M. Heyns, A. W. P. Thomson

outputs to robust operational systems. The SWIMMR program was divided into 11 projects to deliver new capability for forecasting impacts of severe space weather. The Met Office Space Weather Operations Centre (MOSWOC) adopted the new capabilities and products into their existing monitoring and modeling suite.

SWIMMR had a number of high-level objectives including the mitigation of the potential risks of space weather to the UK electrical power distribution network, which is the focus of this work. Our part of the program was named, in an acronym of an acronym, SWIMMR Activities in Ground Effects, or SAGE. A major goal of SAGE was to produce an updated and improved set of software products and services for forecasting the flow of Geomagnetically Induced Currents (GICs) in power grids, high pressure pipelines and railways and make them available through MOSWOC. An additional aim was to improve and operationalize a new set of published scientific models for predicting the ground magnetic field based on solar wind measurements at the L1 Lagrange point.

The SAGE project ran from June 2020 to March 2024. The main tasks were led by British Geological Survey (BGS) who provided nowcasts of GICs in the high voltage power lines, railway lines and pipelines of Great Britain (GB) using real-time minute-mean data from the UK geomagnetic observatories (Beggan, 2015). Researchers at British Antarctic Survey (BAS), Imperial College, London, and University College, London (UCL) implemented a series of previously published scientific algorithms to provide forecasts of magnetic field variation on the ground and contextual information on storm intensity and longevity using machine learning and physics-based magnetohydrodynamic (MHD) models using L1 solar wind measurements (Eggington et al., 2022; Mejnertsen et al., 2018; Shore et al., 2019; Smith et al., 2020, 2021).

The pre-existing MOSWOC capabilities relied on the use of global indicators such as Kp nowcasts (Matzka et al., 2021) and forecasts incorporating impacts of results from the WSA ENLIL Cone model of the solar wind and Coronal Mass Ejections (CMEs) (e.g., Parsons et al., 2011) plus magnetic field observations from three UK sites to infer ground impacts (Clarke et al., 2013). The SAGE project provided new estimates for GICs for the power, gas, and rail systems in the UK based on a new geoelectric field model and new UK-focused magnetic field forecasts based on machine learning and MHD models, all underpinned by new observations (variometers and and magnetotelluric campaign) and augmented by context based models (forecasts of substorm onset, sudden storm commencement and magnetic field rate of change exceedence). This represented an enormous improvement in capabilities.

Several new science models and data sets were created: (a) a new GIC impact model for the high pressure gas pipeline network; (b) a new impact model capturing a railway GIC index; (c) an improved high voltage power grid GIC impact model; (d) an operational version of a global MHD model (GorgonOps) with a novel application to UK magnetic field forecasts; (e) three new magnetic field variometer sites (for improved monitoring and verification); (f) a new magnetotelluric survey of Britain (leading to improved geoelectric field model); and (g) four new machine learning forecasting models of the magnetic field variation and behavior based on L1 solar wind data. We describe these models in more detail in the next sections.

The other novel aspect of the SWIMMR program is the consistent use of cloud computing platforms and modern software tools, notably for version control and containerization, to enable a smoother R2O process. In addition, the required geophysical data and models as measured or created by third parties or partners were made available via a shared Met Office database and user specific web service. The scientific teams were also supported by a dedicated software development team working for the SWIMMR program under direction from Met Office. These initiatives helped overcome the difficulties in traditional approaches of commissioning new software projects that rely only on scientists to move from bare-bones code to professionally developed code with appropriate documentation, logging and error trapping. It was imperative to avoid the usual pitfalls of the "valley of death" encountered when moving from lower tiers of the technology readiness levels to operational systems (Butler & Keller, 2021; Halford et al., 2019).

The use of a cloud based platform and common development tools meant that the project was effectively hardware-agnostic for most services. The set of SAGE models are hosted by the Met Office on their chosen cloud-based platform, presently an Amazon Web Services (AWS) instance. The nowcasting and forecasting outputs consist of 14 operational models, available for project partners and the Met Office on their internal portal, and through the Met Office's Application Programming Interface (API).

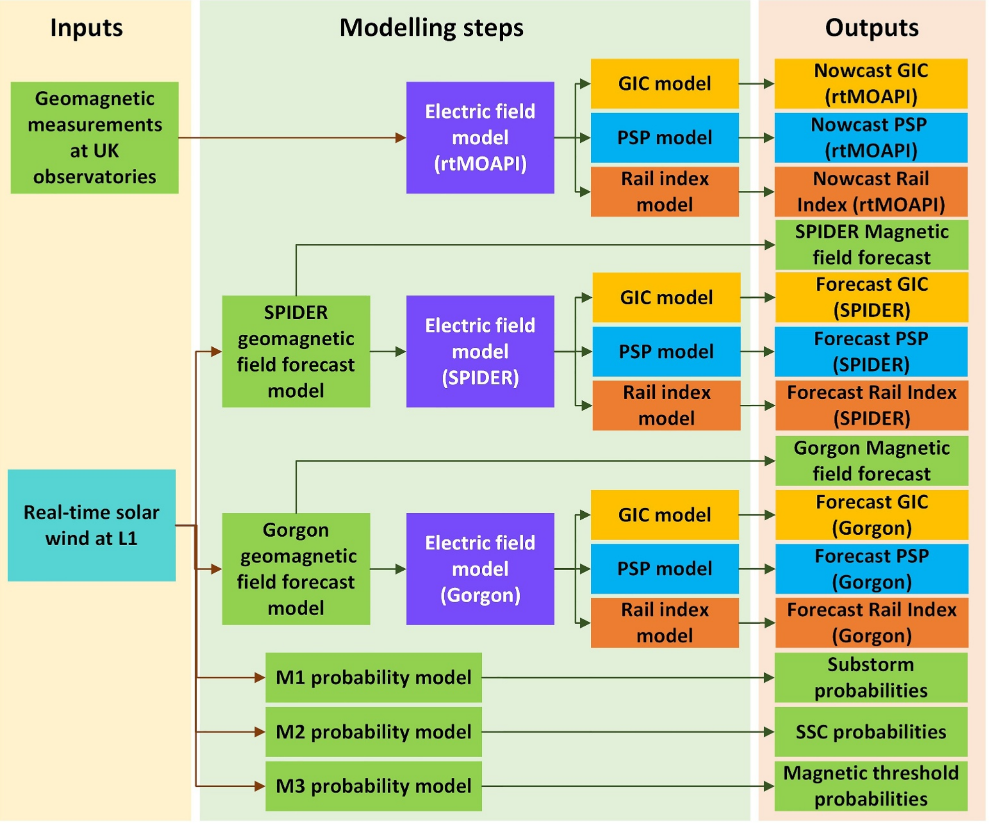

Figure 1. Inputs, processing steps, nowcast and forecasting products for the SAGE project.

In this paper, we outline the scientific basis for each model, and explain the implementation for the R2O process, as this differs somewhat from typical norms for research code. We describe the manner in which data flow through the system, code sequencing and the outputs. We include examples from a G5 geomagnetic storm to demonstrate the utility of each product. Finally, we discuss the lessons learned. A more detailed set of reports on the outputs of the SAGE project can be found in Beggan et al. (2024) and Hübert, Eaton, and Beggan (2024).

2. Model Background and Improvements Within SAGE

The aim of SAGE was to operationalize the UK's space weather nowcasting and forecasting capability for GICs in grounded infrastructure. Prior to the start of the project much of this capability existed as scientific-level research models, developed to investigate a particular question or problem. A primary goal was to convert these models to use standardized inputs and to produce standardized outputs. This was to enable end-to-end coupling of all the models, to allow robust and maintainable nowcasts and forecasts of the environment, subsequent impacts, and forecaster-decision context aids. Figure 1 shows the inputs, data processing steps and outputs for the SAGE project.

Starting with measurements of external magnetic field variation over the UK, or forecasts based on solar wind parameters, the induced geoelectric field can be computed if the subsurface conductivity is known. Using the derived geoelectric field maps, models of GICs in the high voltage power network, rail network or high-pressure gas pipeline network were created at BGS; the latter two models are based on the work of Boteler (2013). These types of infrastructure are electrically connected to the ground either at single points (like electrical substations) or almost continuously in the case of the rail or pipeline network. The conductive metal materials offer a long and low resistance path for the equalization of current across large distances.

The impact of the GIC flowing through each type of the infrastructure varies. In the case of the high voltage power grid, step-up/step-down transformers connected to the ground can experience a bias in the hysteresis cycle with magnetic flux escape from the core and increased reactive power demand causing effects such as even voltage harmonics leading to operational instability (Boteler, 2019). Electrified rail lines can experience signaling faults, which depending on the system configuration can give rise to "right" or "wrong side" failures (Patterson et al., 2024), while pipelines experience excess pipe-to-soil potential which can cause long term damage from increased corrosion rates (Ingham et al., 2022).

The outputs for the GB high voltage grid model have been developed and validated in several studies over recent years using indirect measurements of GIC via the Differential Magnetometer Method (Hübert, Beggan, et al., 2024; Hübert et al., 2020) as well as occasional measurements from Hall probes at four locations in Scotland. The pipe-to-soil potential and railway impact estimates were developed in a previous space weather research project called SWIGS and are novel for the UK so have not yet been validated and are, at present, indicative rather than quantitative. The pipeline and rail models were developed using open source data sets (see Open Research Statement).

2.1. Nowcasting

For nowcasting of GICs, BGS had previously created a real-time magnetic field data collection system at their three UK geomagnetic observatories, and implemented a distribution system via the Internet through portals such as INTERMAGNET and the World Data Centre (Clarke et al., 2013). They had also developed a real-time GIC prediction model for the GB high voltage power network (Kelly et al., 2017) via a web-accessible portal. In the SAGE implementation, the measured magnetic field data from the UK observatories are convolved with a magnetotelluric transfer functions of the subsurface conductivity to generate a map of the estimated geoelectric field experienced on the surface within the past few minutes (Beggan et al., 2021).

The ground electric field is dependent on the electrical conductivity of the subsurface, a parameter that varies widely with lithology type, water content and other rock properties that determine the permeability and porosity. In Britain, the highly variable lithology and its long geologic history cannot be captured in a simple model of the subsurface conductivity. Instead, more complex representations of the subsurface have been constructed for space weather hazard assessment, including a thin-sheet model based on airborne EM data, bedrock sampling, and legacy EM deep-sounding data. However, the thin-sheet model (used as a stopgap during development) has been shown to under-estimate the variability of the geoelectric field. Beggan et al. (2021) demonstrated this at the UK observatory sites where the geoelectric field has been recorded since 2012. Using magnetotelluric (MT) data to derive the ground electric field gives a closer estimate to the measured electric field data. This motivated a new island-wide measurement campaign to improve the conductivity model of Britain (Hübert et al., 2025).

The GIC (in amperes) is estimated in the high voltage grid at each of the nodes in the network. A node is typically a bus or a substation grounding point, though the model also contains features like line splits or unearthed parts of the network (in which case GIC is zero). The present model has over 1,200 nodes with around 1,500 connecting lines. In the online visualization of the outputs (examples of which are shown later), the location and magnitude of GICs in the power grid are provided in an interactive map with a table noting the top five locations with the largest GIC. The map shows the three network voltage levels (as colored lines) along with area-scaled colored circles to indicate GIC magnitude at each substation. The colors (blue, green, orange and red) indicate whether a preset threshold (25, 50 or 75 A) has been crossed. For further context, two smaller plots show the external field variation in the North and East components of the magnetic field at the three UK observatories.

The pipe-to-soil potential (PSP) can be computed using a similar method to modeling the flow of GICs in the high voltage power network with the admittance representation (Boteler, 2013). For the GB high pressure gas network, the pipe-to-soil potential is calculated at thousands of points along the pipelines. The output is an estimate at PSP in volts. For the visualization of pipe-to-soil potential output, a map is drawn indicating the polarity and magnitude of the PSP in volts (V) as color intensity. As with the high voltage power grid, the PSP estimates are based on the largest variation of the magnetic field (and hence geoelectric field) in the past 5 min. Values of PSP outside the range of -0.85 to -1.2 V are considered to be beyond the normal range for safe operation of the pipeline.

The rail index is generated by computing the geoelectric field incident to the direction of the rail track in the 10×10 km grid cell of the geoelectric field map. This value is then scaled to the magnitude of the modeled value during the 13–14 March 1989 storm to produce a value between 0 and 1 (though a larger storm will produce values ≥ 1). This indicates the comparative strength of the geoelectric field and its direction in a dot product form with the orientation of the rail line in the area. The output is a value in each of the 82,000 line sections modeled. Values above 0.8 could indicate the risk of misbehavior of rail side equipment, though further research is needed as this depends on a large number of factors such as equipment type, age and electrification characteristics (Patterson et al., 2023).

2.2. Forecasting

For forecasting magnetic field changes at the Earth's surface using data from spacecraft, the inputs are solar wind measurements from the ACE or DSCOVR satellites at the L1 Lagrange point. Two forecasting models were implemented: (a) the Spatial Information from Distributed Exogenous Regression (SPIDER) model uses machine learning to determine the variation of the magnetic field at ground level in the UK based on training with the largest geomagnetic storms between 1997 and 2016 (Shore et al., 2019); and (b) GorgonOps is a physics-based model of the global magnetosphere based on the Gorgon MHD model (Mejnertsen et al., 2018).

The SPIDER model is a multi-variate regression of the storm-time ground level magnetic field depending on day of year, time of day, and the solar wind epsilon parameter. It relates solar wind parameters to ground magnetic field variation at individual observatories, in this case Lerwick, Eskdalemuir and Hartland. The epsilon parameter captures the solar wind drivers and includes the (absolute) solar wind velocity, the solar wind magnetic field magnitude and clock angle of the Interplanetary Magnetic Field (IMF) vector as it reaches the L1 point. The day of year and time of day parameters are chosen primarily to capture the influence of varying ionospheric conductance from seasonal solar irradiance and particle precipitation. The time of day also accounts for the spatial structure of the ionospheric current with magnetic local time.

The GorgonOps model is an operational configuration of the underlying Gorgon MHD model. The model makes use of an explicit Eulerian formulation of the resistive MHD equations, implemented on 3D staggered Cartesian grid satisfying $\nabla \cdot B = 0$ to machine precision via the vector potential. Various optimizations have been included for operational use cases, including the use of a stretched grid to better resolve the near-Earth inner boundary while improving stability and computational speed. GorgonOps comes pre-initialized with a primed magnetosphere, and appropriate dipole tilt. Coupling to the ionosphere is through a thin shell model mapping to the inner boundary using dipolar field lines, with prescriptions for empirical solar EUV, MHD derived auroral, and background polar cap conductances. Inline solar wind input ingestion, simulation restarts and ground geomagnetic field estimation through line of sight Biot-Savart integration are all included in the operational code.

As well as the required ground geomagnetic field estimation, GorgonOps and its full range of outputs provides additional capabilities via visualization of and information on the wider magnetosphere and the solar wind parameters in near-Earth space. These can be used as contextual information on the severity of a storm due to the proximity of the magnetopause to the Earth or the location of the auroral oval over the UK. Within the context of SAGE, visualizations of the magnetosphere (dynamic pressure) and ionosphere (horizontal current density) are produced every 5 min and all relevant time instances are included for given forecast timestep, that is, the full forecast horizon.

Note that in Figure 1, the code used to produce the geoelectric field for GIC, PSP and Rail index is identical. The outputs differ depending on whether these models are driven by real-time magnetic field observations (rtMOAPI) or magnetic field forecasts derived from solar wind data from either SPIDER or GorgonOps.

2.3. Contextual Products

There are three other machine learning-based contextual forecast models which provide a longer term (hours to days) estimate of how an on-going geomagnetic storm might develop. The first computes the probability of a substorm occurrence (labeled Substorm Forecast) which predicts the next occurrence of the Dungey cycle substorm expansion phase (Maimaiti et al., 2019). Substorms are the indirectly driven portion of geomagnetic storms, which occur when magnetic reconnection in the tail drives energy back into the ionosphere, energizing the auroral oval (Freeman et al., 2019). The second machine learning model is the sudden storm commencement (SSC) occurrence model (known as Shock Impact Assessment) which estimates the probability of an SSC, based on solar wind conditions at L1 (Smith et al., 2020). The third model estimates the probability of extreme rates of change of the horizontal magnetic field (dB_H/dt) occurring at the three UK observatories, called the Extreme

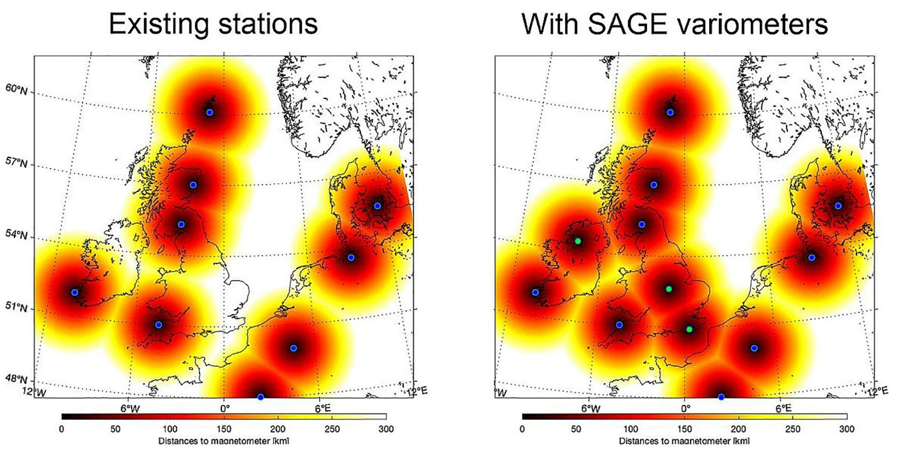

Figure 2. Distance between magnetic sensors around the UK and Ireland in 2022. Blue dots are locations of existing observatories or variometers. (left) Prior to installation of three new SAGE variometers; (right) with the new variometers (green).

Threshold Exceedance Forecast. It computes a series of estimates of the chance of predetermined thresholds (18, 42, 66 and 90 nT/min) of dB/dt being exceeded at each location, based on Smith et al. (2021).

The Substorm Forecast (denoted as M1 on Figure 1) forecasts the occurrence likelihood of a substorm in the next 1 hr period. It is based on a Convolution Neural Network applied to L1 data that predicts substorm onset likelihood (Maimaiti et al., 2019). The Shock Impact Assessment (M2) is an optimized machine learning model applied to L1 solar wind data predicts storm sudden commencement (SSC) within the next 4 hr, based on Smith et al. (2020). SSCs and the 3 days that follow represent over 90% of extreme (greater than the 99.97th percentile) magnetic field fluctuations in the UK (Smith et al., 2019) and more generally below 60° latitude (Smith et al., 2021).

The Shock Impact Assessment (M2) model is split into two submodels: (a) a shock in the solar wind identifier, and (b) an SSC forecast, both using 30 min of data obtained at L1. Statistical features are extracted from 30 min intervals: the mean, minimum, maximum, and range of each parameter (solar wind velocity, density, magnetic field strength and three components of the IMF). These features were scaled (using their mean and standard deviations) and then ranked using their "feature importance" using a random forest model to identify those that are most useful in making the forecasts (as in Smith et al. (2020)). The shock identification model, uses the top three most important features (|B| Range, V Range, Maximum |B|). These are fed to a simple Logistic Regression model that produces a probability that the 30 min interval contains a solar wind shock. If the probability output is greater than 0.7 then it is "accepted" that there has been a shock, and the output of the second model should be consulted. The (b) part of the model uses the top seven most important solar wind features extracted from the 30 min interval (maximum, mean and minimum V, minimum and mean Bz, and the minimum Bx and |B|). The model predicts the probability that an SSC will occur in the next 4 hr using a Gaussian Process model (Smith et al., 2020).

The Extreme Threshold Exceedance Forecast (M3) model is a Convolutional Neural Network applied to 60 min of L1 data that predicts the likelihood of the rate of change of the magnetic field (dB/dt) being above a set of thresholds at each UK observatory, based on Smith et al. (2021). The thresholds are selected from commonly applied values for mid- and high-latitude locations in the literature: 18, 42, 66 and 90 nT/min (Pulkkinen et al., 2013). The probability of exceedance of each threshold is then calculated from the trained relationship to the solar wind parameters for each observatory.

3. Improved Measurement Capabilities

Gannon et al. (2023) argued for the continued support and expansion of ground-based networks to complement and compensate for gaps in satellite data. To this end, part of the SAGE program was to enhance the magnetic field coverage across GB and Northern Ireland. Three new variometers were installed in Florence Court (FLO), Fermanagh, Market Harborough (LEI), Leicestershire and Herstmonceux (HTX), Sussex. Figure 2 shows the effect of three new variometer sites in reducing the distance to the nearest magnetic sensor to below 200 km across the land areas.

The variometers record data at 1 Hz cadence and provide real-time data back to the BGS office every 5 min via the internet or mobile phone network. Each system uses a Sensys 3-axis fluxgate magnetometer connected to an EarthData 24-bit digitizer (Hübert et al., 2020). The FLO and LEI systems are completely isolated and self-contained, sitting in open farmland away from anthropogenic noise sources. At each site, a 270 W solar panel connected to two 95 Ah batteries provides power for the electronics and a 4G/LTE modem; these draw around 10 W on average which is easily recharged daily by the solar panel even during winter. Occasional site visits are necessary to check the status of the equipment. The HTX sensor is sited within the grounds of the Herstmonceux Satellite Laser Ranging Facility and is powered from the electrical mains and connected via wired Ethernet to the BGS network.

The sensors are installed and stabilized on wooden mounts in shallow pits in waterproof containers and connected to the control electronics around 6 m away to reduce magnetic interference. They are oriented toward magnetic north and east by precise leveling of the Z axis and then nulling the Y axis to 0 nT on the day of installation. Thus the sensor X axis points along the H component of the magnetic field. As the declination angle is presently close to zero in Britain, the misalignment between the H and X components at each site is small. The three systems were installed between February and July 2022.

We estimate the nominal measurement noise level is 0.1 nT, determined from calibration of each sensor at Eskdalemuir observatory. While FLO and LEI are almost entirely free from man-made interference, HTX experiences almost continuous occurrences of spikes and fluctuations of ±5 nT from local electrified train line ground loops. The noise vanished entirely, for example, during a series of rail strikes in 2023. In large geomagnetic storms, the effect from the rail noise becomes negligible (<1%).

As noted in the Introduction, the response of the ground geoelectric field to magnetic field variations was not well known. Between June 2020 and March 2024, an extensive field campaign was undertaken to measure ground electric and magnetic field variations at 53 locations across mainland Britain using the magnetotelluric (MT) method to characterize the local geoelectric response during increased geomagnetic activity. Additionally, we made use of another 17 MT measurements captured in the UK between 1990 and 2020. Hübert, Eaton, and Beggan (2024) provides further details of the deployment campaign, instrumentation and processing. MT measurement sites are shown in Figure 3. These measurements allow us to model the space weather response of the subsurface using MT transfer functions as well as image deep geological structures by inferring the subsurface conductivity through 3D inversion. Operationally, the new MT transfer functions drive the geoelectric field computation, which replaced the original thin-sheet modeling code used as a placeholder while the MT and code development ran in parallel. We use the methods of Campanya et al. (2019) and Malone-Leigh et al. (2023) to implement a magnetotelluric-based model to produce a geoelectric field maps of Britain. Hübert, Beggan, et al. (2024) demonstrated the improvements using MT data set over the thin-sheet modeling approach for GIC calculations.

4. Implementation on a Cloud-Based Platform

The SAGE applications are written in Python and run on Met Office-provisioned infrastructure. The compute-heavy GorgonOps modeling application runs on an on-premises High Performance Computer (HPC) while everything else runs in the cloud, presently on AWS. The architecture and software development workflow is comparatively private and stringent by academic standards. This is driven mainly by Met Office requirements for the security of their compute platforms, and to improve the operational readiness of the models by requiring operational-like development workflows early in the R2O process.

An advantage of this approach is to allow long-term attribution of changes made to the model code or environment to the resulting effects on model output. The development workflow used a relatively small amount of overarching software: Docker, git, and GitHub. Docker is a software application that "containerizes" the code, data and underpinning software for each space weather application into a self-contained and executable image or virtual machine. When run within a cloud computing environment, each separate Docker application is capable of acquiring and processing inputs (such as magnetic field and L1 solar wind data), and outputting the relevant products. The advantage is that any problem or crash does not affect the overall system and the affected container can be restarted. Another feature called Docker-Compose is used to handle orchestration between containers for each application so that intermediate data products (such as the geoelectric field maps) can be passed between them.

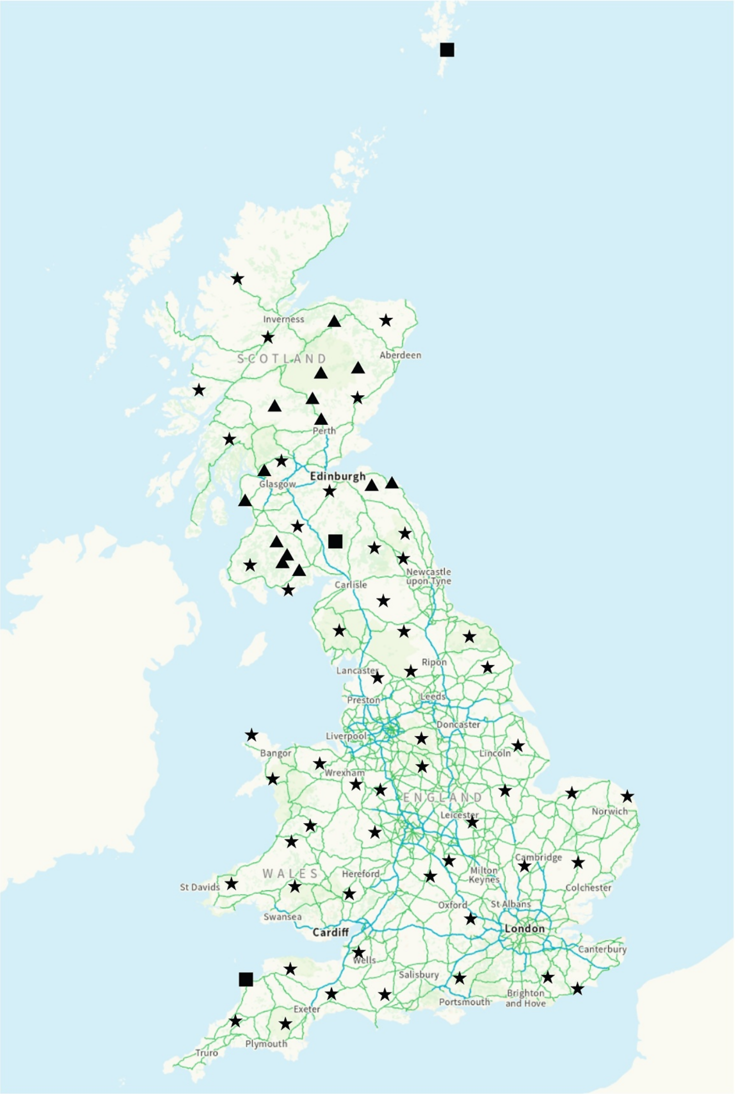

Figure 3. Location of magnetotelluric survey sites across Britain and geomagnetic observatories (squares). 53 (indicated with *) sites were collected during the SAGE project. Legacy data from previous projects are indicated with triangles. Green lines are the major road network with blue lines showing the motorways.

The git application is used for version control of the resulting Docker code, which together define the overall *deployed* version of each SAGE model. The GitHub website is used to host the SAGE git repository (or "repo") which acts as a target for SAGE developers to push the developed application code and subsequent changes to, and as the source from which to deploy updated versions of the applications onto the AWS environment. The entire develop and deploy process is automated so that a git push of new code to the main branch of the repo triggers the full application build, upload and execution on the cloud platform.

Given this role, for operational security, the repo is configured to be private with only SAGE partners and Met Office staff having access. Furthermore, reproducibility and security is enforced via configuration restrictions on AWS access: changes to the deployed applications running on AWS can only be made by pushing changes to the GitHub repository, thus leaving a permanent log of the change made and allowing unambiguous attribution. However, to avoid this restriction from inhibiting initial development workflow, the containerized applications can also be run locally on partner infrastructure for coding and testing. This approach of running on equivalent Dockerized setups on partner computers, prior to pushing to GitHub and deploying to AWS, allows faster development, creation of new features, testing of changes, and simplifying the process for debugging any issues

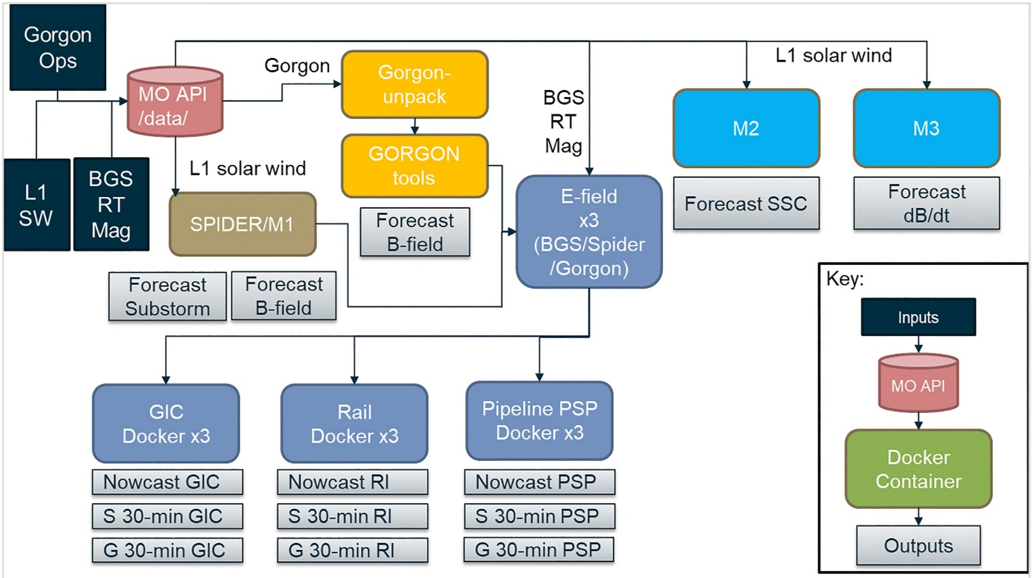

Figure 4. Docker implementation for the SAGE project.

arising. The resulting Docker containers work in tandem to produce the now- and forecast GIC data streams. Figure 4 shows an overview of how the Dockers on the cloud computing platform receive and pass data between the various parts and process the inputs. We next describe the data inputs, flow and outputs in more detail.

Development and documentation of the evolution of the modeling code from scientific outputs to operational software was captured using the concept of Application Usability Levels (AULs) as described by Halford et al. (2019). AULs are similar to the idea of a Technology Readiness Level (TRL) applied within engineering design. A TRL has a value of between 1 and 9, where 1 is a very early conceptual model and 9 is a robust real-world system. For AULs, Level 1 equates to basic research through to Level 9 which is approved for use in an operational environment.

All SWIMMR models were deemed to have completed initial stages (basic research, establishment of requirements, and assessment of state of the art) via prior work and a successful SWIMMR grant proposal, and were defined to be at AUL3 at outset. The initial model codes for SAGE were judged to be of AUL4 (initial integration) and were raised over time to AUL6 (complete validation) via internal development at each partner. The next stage was to implement the codes on the cloud-based platform (AWS) to make an application prototype. The SAGE partners took the system to AUL8 demonstrating validation in relevant “real world” environment. Full implementation to AUL9 (approved for on demand use, in operations) rests with the Met Office as they integrate it into the MOSWOC suite.

4.1. Accessing Data Through the Met Office API

The Met Office provide a dedicated API allowing only specific magnetic and solar wind data to enter the SAGE Dockers. The API is an access point for a MongoDB database which is controlled through calls to a password-key protected set of web links (often termed endpoints). Endpoint web addresses are followed by text linking to the specific data product (e.g., the magnetic field at Lerwick observatory or the solar wind velocity). These can be programmatically accessed using standard internet protocols in Python.

To fulfill the data requests, the Met Office ingests magnetic and solar wind data from three sources. The first is the real-time solar wind data from the NOAA Space Weather Prediction Centre (SWPC). These data are collected by a number of spacecraft, specifically ACE and DSCOVR, and include the solar wind velocity, density, IMF magnitude and orientation and the particle temperature at the L1 Lagrange point. There are occasional gaps in this data set when the L1 satellites are not able to downlink data, for example, leading to a delay, or if instruments are offline for maintenance which can have an impact on the forecasting products (Smith et al., 2022).

The second data set is real-time magnetic measurements to the API from the three UK observatories (Lerwick, Eskdalemuir and Hartland) in two formats—as full field values in the three linear orthogonal components (X, Y, Z) and as external field only with core and crustal contributions removed. The data are provided by BGS every minute and ingested into the Met Office database.

The third data set is ionospheric, magnetospheric, and ground magnetic field estimates generated by the GorgonOps model. GorgonOps is configured to run continuously in real-time on the internal Met Office HPC environment, itself ingesting real-time solar wind data from the API. The GorgonOps operational model has an inner boundary radius of 3.5 Earth radii (R_E). This model computes plasma quantities of interest throughout the simulation domain on a stretched grid, which spans from 30 R_E upstream to 130 R_E downtail, and an extent of (−60, +60) R_E in the other components. In addition, relevant ionospheric and magnetospheric parameters are computed, along with estimates of the ground geomagnetic field variation (Eggington et al., 2022).

When the GorgonOps model is spun up, forecast data from the HPC take a few iterations to become available, so the ground geomagnetic field estimate is given a status flag. In the case of missing data due to spin-up or restart, a placeholder is used to denote this. Solar wind data are ingested as is, with the forecast profile shifted as appropriate relative to the solar wind speed. This adaptive forecast regime allows the model to be as robust and accurate as possible, with the latest information received from L1 being the most relevant so a full forecast profile is maintained. This is important when shocks or more rapid solar wind conditions arrive. Once the model is running, its outputs are packed and transferred back to the API every 5 min.

4.2. Magnetic to Geoelectric Field Maps

On the AWS cloud computing environment, the GorgonOps data "unpack" Docker runs every 5 min, extracting the relevant files computed on the HPC for a given forecast horizon determined by the current solar wind speed. The resulting magnetic time-series is shifted relative to the forecast horizon. Measured ground geomagnetic field data from the UK observatories are used to determine the baseline level of the variation. Together with the real-time magnetic data this provides a forecast of magnetic field variation up to 1 hr ahead.

The SPIDER model forecasts magnetic field variation at LER, ESK and HAD between 30 and 60 min ahead of time, depending on the velocity of the solar wind at the L1 Lagrange point. These measurements are provided by the Met Office API and the machine learning model forecasts the external magnetic field variation using 2 hr of preceding measurements.

The output from SPIDER and GorgonOps can be treated in the same way as the real-time magnetic data, to estimate a geoelectric field map across Britain, and hence GICs in the power, pipeline and rail networks.

The geoelectric field Docker reads a magnetic field time-series from either the real-time BGS external field measurements or with the concatenation of the measurements with the forecast from SPIDER or GorgonOps. The magnetic field time-series are convolved with each of the 70 magnetotelluric transfer functions (MTTFs). The MTTFs use magnetic field time-series over the past 48 hr collected from the API. The time-series is then padded with another 105 min of magnetic values tapering to zero, or is shifted to length of the forecast from SPIDER or GorgonOps, and further extended to 105 min tapering to zero. The magnetic field time-series values are computed from the latitude-weighted measurement using the two closest observatories. Once the geoelectric field time-series has been computed at each of the 70 MT sites, the minute containing the largest geoelectric field value in any of the sites in the final 5 min of the real-time (or forecast period) is selected. The geoelectric field values, for this minute, are spatially extrapolated across Britain using a cubic spline and are sub-sampled to 10 × 10 km cells.

The electric field container is replicated three times in the AWS and generates a snapshot map of the geoelectric field across the UK mainland at a 10 km spatial resolution for the present period if using the BGS observatory data, or for a forecast point forward in time if using SPIDER or GorgonOps. The map (either nowcast or forecast) is passed internally within the Docker environment to the GIC Dockers to compute the values for the high-voltage power grid, railway index and high-pressure pipeline.

4.3. GIC Products

There are three outputs from the GIC-specific Docker containers which are: (a) computed GIC in the high voltage transformers on the 400, 275, and 132 kV (Scotland-only) transmission system, (b) the Pipe-to-Soil Potential (PSP) in the high-pressure gas pipeline network and (c) a rail index indicating where potential issues in rail signaling might occur in the GB network. These accept the geoelectric field map as input, and output data files and visualizations showing the location and magnitude of the GICs.

The three types of input magnetic fields (real-time measurements, SPIDER and GorgonOps) are replicated three times, giving nine separate Dockers to output nine GIC products. In quiet time conditions, the nowcast and SPIDER operate continuously. When storm conditions are forecast, the GorgonOps model is spun up. As it is very computationally intensive, running it continuously would be an inefficient use of HPC resources.

4.4. Contextual Forecast Information

There are three additional contextual forecasts named the Substorm Forecast (labeled M1 in Figure 4), the Shock Impact Assessment (M2) and the Extreme Threshold Exceedance Forecast (M3), all of which are machine learning models. These predict (a) the probability of a magnetic substorm occurring, (b) the probability of a storm sudden commencement occurring and (c), the probability of the rate of change of the horizontal magnetic field exceeding a set of threshold levels at each of the UK observatories in the next few hours. These models are driven by solar wind data from the L1 point. They output files containing time-series of the probability (between 0 and 1) of the likelihood of occurrence of each parameter.

5. Operational Sequencing

The input data within the API are updated every minute. Most of the SAGE models require at least 5 min to several hours of data in order to produce an output. There is also a trade-off with the required processing time and useful update rate of the model. Processing time is on the order of 1 to 2 min, depending on the model, so that is the reasonable lower limit. At the upper end, updates of 30 min would likely be too long for operational utility.

The system is therefore set, at present, to update all products where the appropriate input data are available every 5 min. The processing is controlled by the Linux *crontab* application within each geoelectric field Docker container and set to run the main processing code at 5 min intervals. Once the geoelectric field maps have been created for that 5 min period, "watchdog" code in the other Dockers check if the map has been updated. When this condition is met, the other containers begin their own processing. The outputs are copied out of the AWS to the Met Office database via the API.

As an example, Figure 5 shows the sequencing of the steps for the nowcast part of the code. The geoelectric container (E-field Docker) in nowcast mode triggers every 5 min, reading the most recent BGS observatory data from the Met Office API. It computes a geoelectric field map, creating a *success* file once complete. This ensures that the map output has finished writing to file before any further processing occurs. The GIC, PSP and Rail dockers check every 1 min to ascertain whether the *success* file's timestamp has changed. When it does, this triggers them to compute their respective outputs which are collected when complete by the Met Office API from the */output/* folder.

In the forecast mode, the SPIDER model is run every 5 min, again controlled by the *crontab* functionality in the SPIDER docker. This generates a new file with the magnetic field forecast for the following hour (depending on solar wind velocity). The E-field Docker (in forecast mode) is triggered every 5 min and reads the SPIDER data file if a new file is detected. As with the nowcast sequence, the geoelectric field code runs and creates a *success* file once complete to ensure that the geoelectric field map has finished writing to file. The GIC, PSP and Rail Dockers check every minute to ascertain whether the *success* file's timestamp has changed. When it does, this triggers them to compute their respective forecast outputs for collection by the Met Office API. Similarly, the Shock Impact Assessment and Extreme Threshold Exceedance Forecast models are triggered every 5 min.

For the GorgonOps model, due to the larger and more intense computational requirements, this runs on the Met Office Monsoon HPC infrastructure. When manually activated for periods of high geomagnetic activity, the model produces estimates of the geomagnetic field at a 1 min cadence, along with contextual visualizations. The ingestion cycle for real-time solar wind input on Monsoon is every 5 min, which is linked to a matched file

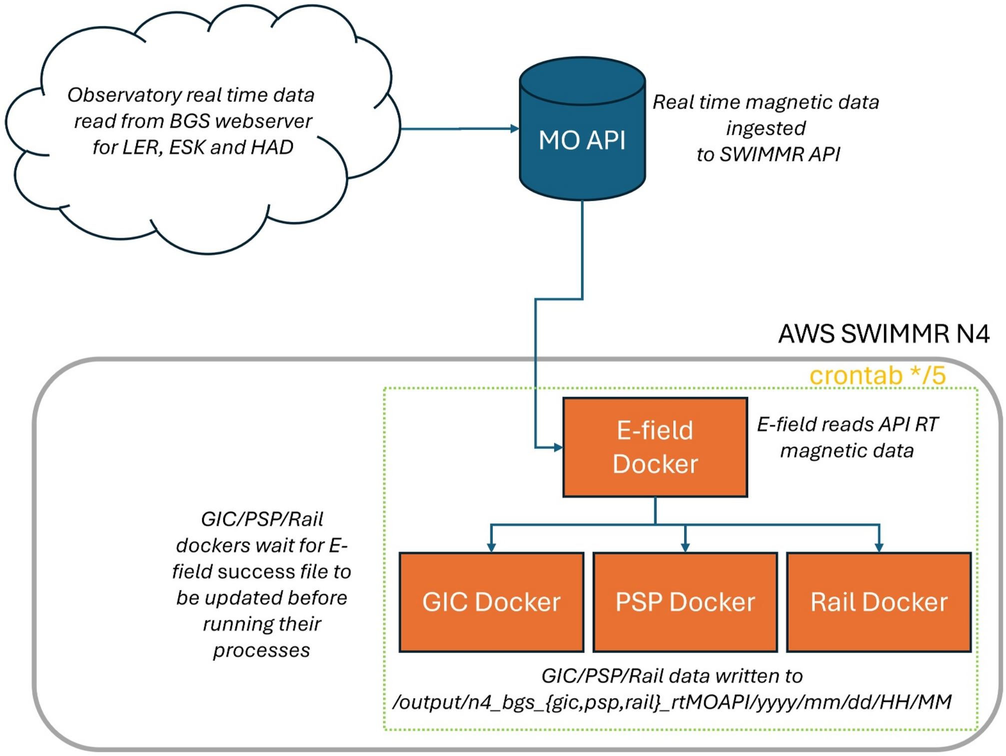

Figure 5. Nowcast sequencing in the SAGE Docker.

transfer cadence. All processing in these cases is greedy, making use of all available data for processing at the point of being triggered. Each iteration generates a new file with the magnetic field forecast dependent on the current solar wind speed for the following hour (with appropriate padding), and is stored in the appropriate directory structure. As with the nowcast and SPIDER, the E-field code in forecast mode is triggered every 5 min and reads the available GorgonOps magnetic field forecast data files, computes the geoelectric field map and creates a *success* file once complete. The GIC, PSP and Rail Dockers check to ascertain whether the *success* file's timestamp has changed, and triggers them to compute their respective outputs.

6. Model Validation

Validation of operational models is important for monitoring behavior, understanding limitations and to establish a baseline for longer-term improvements. It also determines whether AUL9 has been achieved.

Validating the SAGE models is challenging as the "ground truth" for many is not readily accessible. The products thus fall into four general categories in terms of ability to validate them.

- "Readily verifiable": These products can be verified directly, continuously, and automatically in the AWS. The magnetic field forecasts from SPIDER and GorgonOps can be compared to the ground observatory measurements, typically within an hour.
- "Possible offline": These products cannot be verified within the AWS, as post-event data from an external third-party provider are required (e.g., GIC measured in the high voltage power grid) or expert judgment such as whether a substorm or sudden storm commencement has occurred.
- "Measured but not available": Pipe-to-soil potentials in high pressure pipelines are measured for industrial monitoring or safety purposes but are not available for research purposes.
- "Not physical": The rail index is not a physical property so cannot be directly validated.

GIC measurements are presently made at two substations in Scotland at Torness and Strathaven (Thomson et al., 2005). To validate the GIC in the high voltage network model elsewhere in the GB grid, a large-scale campaign ran between 2017 and 2021 to make indirect measurements of the GIC flowing in power lines using the differential magnetometer method. The results from 12 sites in Scotland, Wales and England show excellent correlation and matched the amplitude of the line GIC when the geoelectric field is correctly estimated using MT-derived values (Hübert, Beggan, et al., 2024; Hübert et al., 2020). This give us confidence that the network model is correct, given an accurate geoelectric field model.

While there is ample evidence of GIC flow in high pressure pipelines from studies in Alaska, Finland and New Zealand (Campbell, 1980; Ingham & Rodger, 2018; Pulkkinen et al., 2001), measurements of the PSP in the UK pipeline network are not available to researchers and no proxy measurements of PSP have been made either. The values of PSP predicted are thus theoretical but they are intended to be indicative of where problems may be experienced by operators.

There is strong evidence of GIC effects in rail from measurements in rail lines particularly at high latitudes in Russia (e.g., Eroshenko et al., 2010). The GIC can be modeled using a similar method to modeling the flow of GIC in the high voltage power network with the admittance representation (Boteler, 2013). However, like measurements of the PSP in the UK gas pipeline network, measurements of GIC in rails are not available, as it is a difficult measurement to make on live rail networks (primarily for safety reasons). Future opportunities to improve the modeled effects on railways may arise in future based on new work (e.g., Patterson et al., 2023).

For the contextual forecasting models, these can only be post-validated as they rely on expert assessment made days to months after the event. It is not possible to run any validation on the active, operational system. This situation is less than ideal as data need to be stored for long periods of time, and often other data sets are required. For these reasons validation must be completed offline.

7. Application to the 10–12 May 2024 Storm

Between 3 and 10 May 2024, a series of X-class flares emanated from sunspot region AR3664. Associated with these were between five and seven coronal mass ejections (CMEs). Several CMEs appear to have merged en-route to the Earth, resulting in a complex and long-lasting magnetic cloud structure. The CMEs triggered a geomagnetic storm categorized as extreme (G5) on the 10th May 2024 with over 24 hr of continuous Kp8 and Kp9 magnetic variations recorded (e.g., Díaz, 2024). Variously termed the "Gannon" or "Mothers Day" storm, it ranks at present as the third largest storm in the *aa** index in terms of longevity but not necessarily as an extreme storm by Dst magnitude or other rankings (Elvidge & Themens, 2025).

Around 17:00 UT on 10 May 2024, the absolute IMF magnitude increased from a few nanoTesla to a peak of around 70 nT at 23:00 UT. The solar wind IMF remained significantly elevated throughout 11 May, and did not decrease until early on 12 May. During most of this time, Bz was predominantly southward (negative), driving strong geomagnetic activity across the globe accompanied by vivid auroral displays at mid-latitudes (e.g., Grandin et al., 2024). The peak of the geomagnetic storm in the UK was around 22:35–22:45 UT.

The SAGE AWS system had been implemented in stages, with the nowcast system set up initially, followed by the SPIDER forecast and GorgonOps forecasting capabilities. The machine learning context models were completed in February 2024 making the entire system fully operational by early March 2024. The May 2024 was therefore an excellent test of its utility. The GorgonOps model was manually initiated on the afternoon of the 10th, spinning up from a "cold" start over 2 hr.

We show a series of plots from the operational SAGE system running on the cloud computing service on the 10th May. Figure 6 shows the estimated GIC in the GB high voltage grid. The upper panel shows a snapshot of the GIC in the 400, 275 and 132 kV network with circle size and color indicating magnitude of GIC flow. A table in the upper right shows the modeled GIC at the top five sites. At that moment, Pembroke (PEMB, in southwest Wales) had a modeled instantaneous substation GIC at 138.99 A. The other top sites are NORM (Norwich Mains, eastern England), TRAW (Trawsfynydd) and WYLF (Wyfla) both in northwest Wales and SIZE (Sizewell, eastern England). The large modeled GIC in the middle to south of Britain were relate to the location of the auroral oval which was overhead in central England at this point. These sites are also at the end of long lines at corner nodes in the grid (Beggan, 2015).

For reference, the circles in the visualization are color-coded by an arbitrary set of thresholds of 5, 25 and 50 A. As these are total GIC at each substation, the per-transformer GIC is divided by the number of transformers on site and will be much lower. We also emphasize that the initial GIC estimates were later revised post-event using a

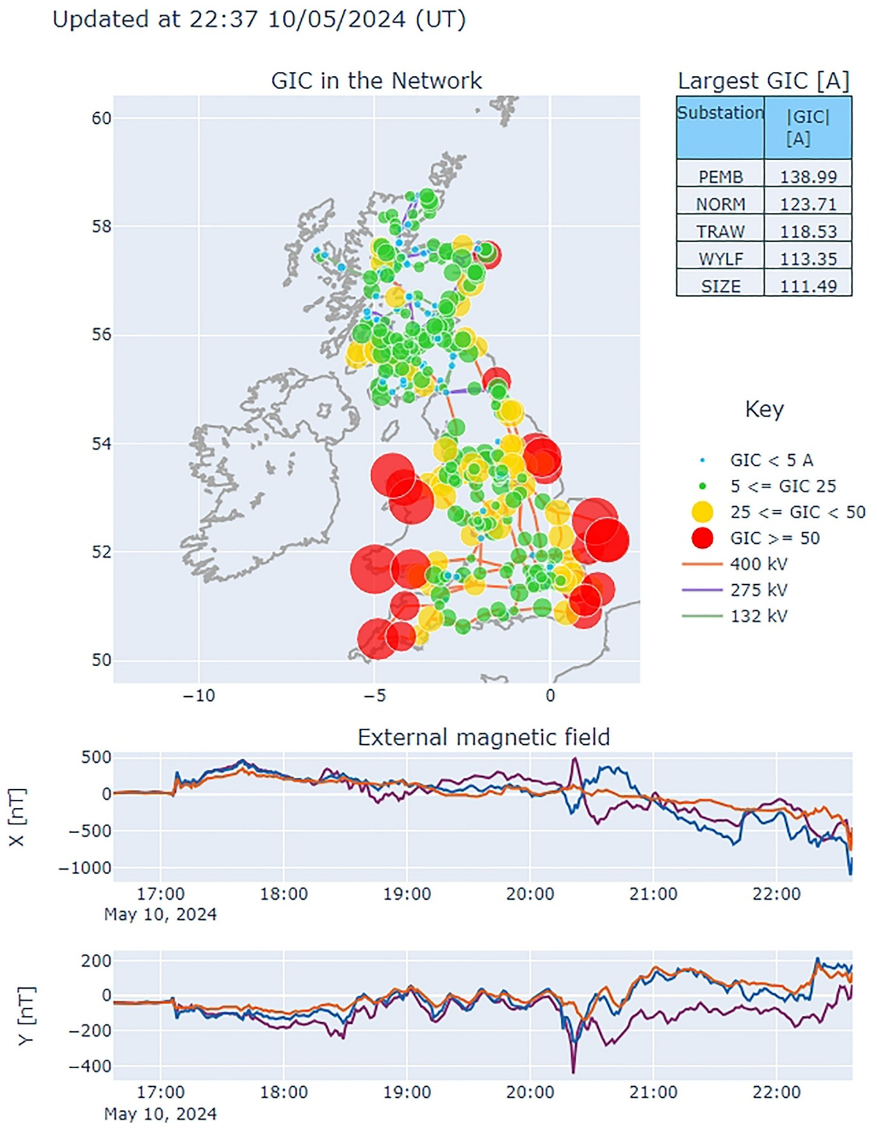

Figure 6. Map of the nowcast of geomagnetically induced current (GIC) (in amperes) in the Great Britain high voltage grid at 22:37 UT on 10-May-2024. Upper panel: GIC magnitude plotted on a map. Circles indicate magnitude at each substation. Line connections are color-coded by voltage level; lower panels: external magnetic field variation (in nT) at LER (purple), ESK (orange) and HAD (light blue) observatories in the North (X) and East (Y) components.

larger number of magnetic observatories and variometers as well as an improved version of the geoelectric field model, which strongly reduced the peak values of GIC (Lawrence et al., 2025). The plots shown here were the output of the system at the time of the storm; in November 2024 the geoelectric field model was upgraded from the thin sheet to the MTTF version. Given the estimated 1-in-30 years nature of the event (rather than a 1-in-100 or greater), there were no reports of mis-operation in the GB power network. The dynamic configuration of the electricity network changes continuously, so some of the locations estimated to have high GIC may not necessarily have been in service. The model is therefore indicative of where high GIC flow is potentially occurring. As with all R2O projects, further feedback from the operations side allows new research to take place and improvements to the model and code to be made.

In the lower panels, the external magnetic field recorded at the three UK observatories are shown in the X and Y component for Lerwick (purple), Eskdalemuir (orange) and Hartland (light blue). The sudden storm commencement at 17:15 is visible. This plot is a snapshot taken from an interactive html-rendered map. For a

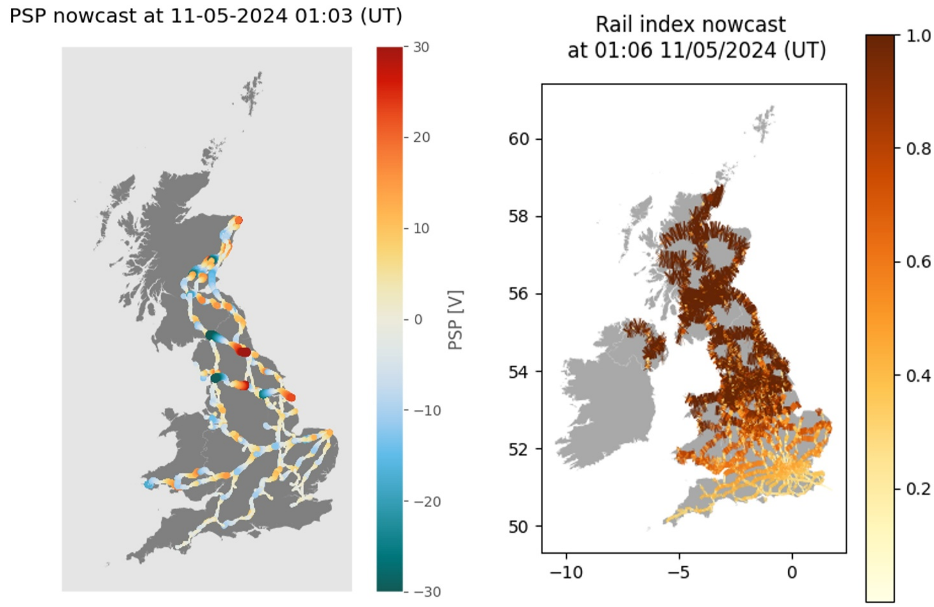

Figure 7. Nowcast of (left) the pipe-to-soil potential (Volts) in the high pressure gas network at 01:03 UT; and (right) the rail index (unit-less) at 01:06 UT on 11-May-2024.

user, it is possible to view the maps in a web browser and zoom in on a particular region and hover over a site with a mouse cursor to query the node name and GIC value.

Figure 7 illustrates a similar set of maps of the estimated pipe-to-soil potential (left panel) and the normalized rail index (right panel) at 01:05 UT on the 11-May-2024. The pipe-to-soil potential (in volts) is indicated by the color and thickness of the lines. Large values occur at the end of longer pipelines where the PSP accumulates. Values of up to 30 V are modeled, which is well outside the ±1 V range required for safe operation. The technique for pipeline GIC modeling remains in development so the values are indicative of where operational issues could occur in the network.

The rail index shows regions of the rail network where the geoelectric field aligns strongly with the direction of the railway line (within a 10 km cell) and the magnitude of the cross product, scaled relative to the maximum modeled values during the March 1989 storm. At the snapshot time, the index suggests northern England and Scotland were experiencing conditions similar to the March 1989 storm. The index intensity over Scotland, north to central England and Wales relates strongly to the auroral oval extent and to the normalization against the peak values modeled in the March 1989 storm. There were no incidents reported on the rail network, though at 02:00 local time, very few trains would be in operation. Signaling failures are expected to be ephemeral during a geomagnetic storm so often may not be observed and the system quickly resets to its normal state (Patterson et al., 2023).

The forecast magnetic field variation and the estimated GIC in the high voltage grid from the SPIDER and GorgonOps models are shown in Figure 8 (left and center) for around 17:55 UT on 10-May-2024. The upper panels with the maps show the largest estimated GIC in the next 30 min (depending on solar wind speed at the L1 location). The lower panels show the measurements of the field at the three UK observatories up to the black vertical dotted line which is the forecast starting point. The external magnetic field recorded at the three UK observatories are shown for the X and Y component at Lerwick (dark blue), Eskdalemuir (orange) and Hartland (light blue). Forward in time shows the estimated variation of the magnetic field in a lighter shade. The forecast times are slightly out of synchronization as the computed horizons of the SPIDER and GorgonOps models are different.

For the SPIDER model (left panels), the solar wind values for the past 6 hr contribute to the estimate of future variation. Three lighter colored lines on these plots indicate the past and future magnetic field variation based on the solar wind. The left panel shows the SPIDER forecast made at 17:58 UT and suggests the magnetic field

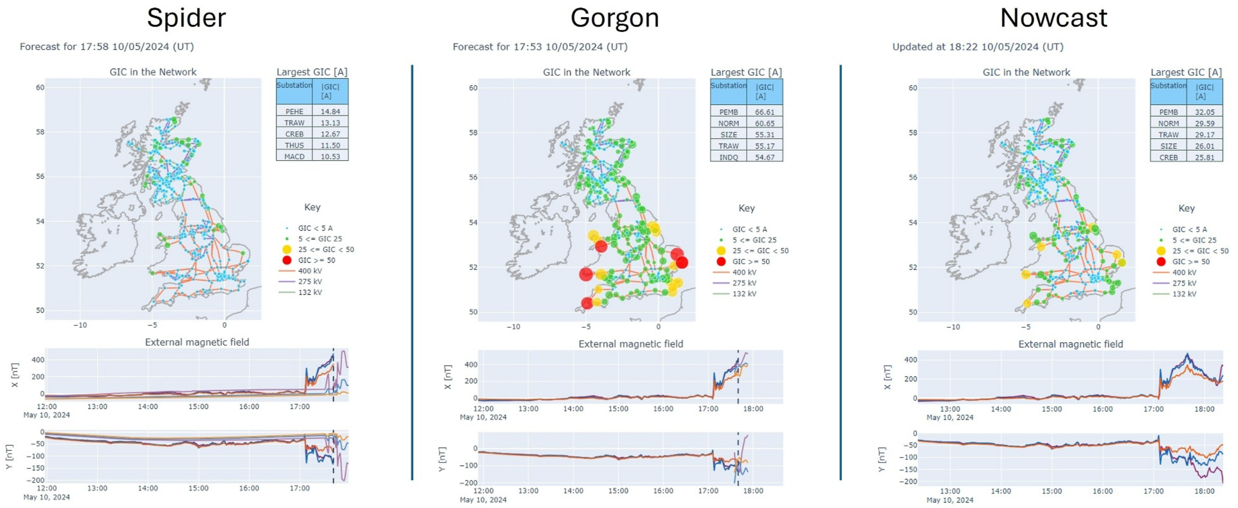

Figure 8. Forecast of geomagnetically induced current (GIC) from spatial information from distributed exogenous regression (SPIDER) (at 17:58 UT) and GorgonOps (at 17:53 UT) on 10-May-2024 compared to the nowcast from British Geological Survey observatory data approximately 30 min later at 18:22 UT. Map of the nowcast of GIC (in amperes) in the Great Britain high voltage grid at 22:37 UT on 10-May-2024. Upper panel: GIC magnitude plotted on a map. Circles indicate magnitude at each substation. Line connections are color-coded by voltage level; lower panels: external magnetic field variation (in nT) at LER (purple), ESK (orange) and HAD (light blue) observatories in the North (X) and East (Y) components. Lower panels show the measured external magnetic field up to the black vertical dashed line, after which the forecast variation is plotted. For SPIDER, additional lighter colored lines show the machine learning variation based on solar wind parameters. See text for details.

variation will decrease, then increase at Lerwick and Eskdalemuir but with a smaller change at Hartland. This creates a relative low estimate of GIC in the forecast period.

The GorgonOps forecast (center panels) was made at 17:53 UT and suggests a continued increase in the external magnetic field variation in the UK, particularly at Eskdalemuir and Hartland. This forecast produces larger values of GIC estimated in central Britain, shown by the appearance of red circles at the coastal edges.

The right panels show the magnitude of the GIC from the nowcast model produced at 18:22 UT, at the approximate time when the SPIDER and GorgonOps models in the other panels are providing esimates for. The nowcast GIC magnitudes lie between the two forecasts made 30 min earlier.

Figure 9 shows the context-based models driven by the L1 solar wind data. The machine learning contextual based outputs provide estimates of the storm likelihood of a shock being detected in the solar wind and whether a storm sudden commencement is expected. Panel (a) shows the solar wind velocity and magnetic field properties at 17:45 UT on 10-May-2024. At the onset of the "Gannon" storm, the shock in the solar wind is picked up in the bottom panel, though as some points are missing in the real-time L1 data set, it makes it difficult to ascribe a high likelihood to the shock. This means the probability of a shock is around 0.5, though the probability of a sudden storm commencement is around 0.9. In panel (b) the likelihood of the five different threshold rates-of-change are provided. These are visualized as concentric rings with the lowest rate of change (18 nT/min) almost certain (>98%) to be exceeded in LER and ESK, but the highest rate of change (90 nT/min) having a 14% chance of being exceeded in LER during the next 4 hr. At lower latitudes of ESK and HAD, higher rates of change are assessed to be less likely.

8. Advances and Challenges

The May 2024 storm was the largest in the past 30 years. It offered an excellent test of the whole system, proving it could respond as expected and to provide useful real-time information. However, the output is, naturally, only as good as the available inputs. Given the difficulty in operating a continuous service of solar wind measurement at L1, there are inevitably gaps in the input solar wind data streams and subsequently in the outputs. The data from BGS-operated geomagnetic observatories are more robust with several layers of redundancy in each step of the processing chain from measurement to delivery. This includes multiple sensors at each site, wired and satellite

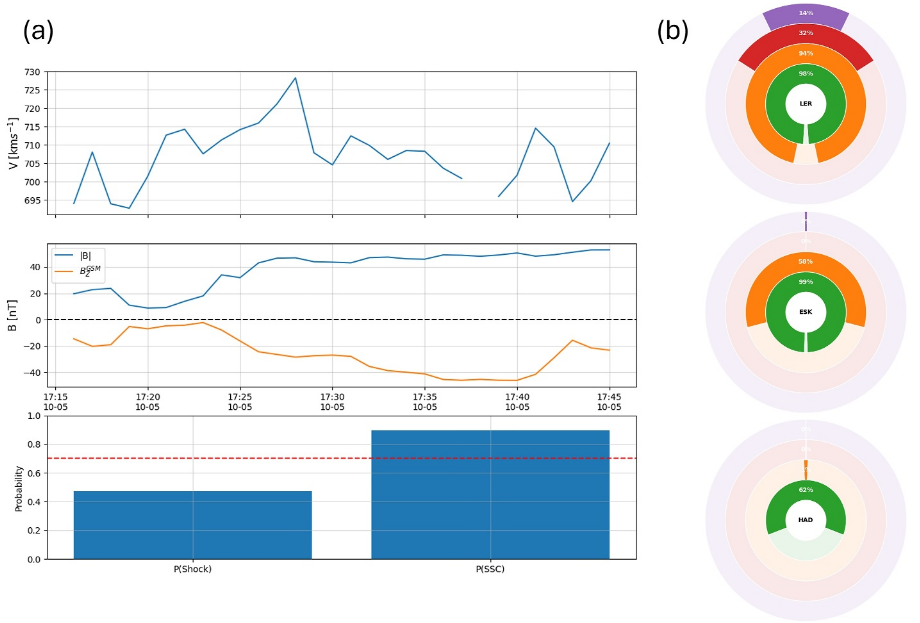

Figure 9. Contextual information from the (a) Sudden Storm Commencement and (b) Threshold Exceedance at UK observatories at 17:45 UT 10-May-2024.

internet connections, triple redundancy processing systems and three independent webservers to output the magnetic field values every minute.

Smith et al. (2022) undertook a thorough study of the availability and quality of L1 data from the scientific releases from ACE and DSCOVR as compared to the available data in real-time (within 1–5 min of measurement). They showed that small short gaps are common in the operational data, and that interpolation for up to 5 min would strongly improve the up-time of the models that require long continuous time-series to operate. A simple solution to interpolate over short gaps of up to 15 min can extend the availability of longer periods of data windows. For example, the SPIDER model requires at least 120 min of continuous data to operate. If no interpolation is applied, then using real-time L1 data could only provide forecasts less than 1% of the time. A small amount of interpolation (e.g., 5 min) raises this to 75%, according to the Smith et al. (2022) analysis. The additional point to make is that many machine learning models are trained on scientific-level data sets (e.g., OMNI) which have been cleaned and calibrated over time. Near-real-time data is often much less well-behaved.

The SAGE project offers an example of the process of research to operations, moving from AUL4 to AUL8 over the course of 3 years. The academic project team worked closely with the MOSWOC science team throughout, with bi-monthly meetings to discuss progress and ask for recommendations for model operation and implementation. The initial design evolved over the project to take account of operational restrictions or resources. For example, GorgonOps was intended to run on AWS but the computational requirements meant it was moved to the Met Office HPC suite with other changes needed for the API inputs and outputs.

Some of the wider lessons learned are that the process of moving from academic to operational code does require professional IT developers working in conjunction with scientists. The learning curve for many of the standard tools (Docker, git) is relatively steep and best-practice does need to be shared. The SAGE project had six partners (four academic plus Met Office and their third-party developer) who worked effectively together with project meetings every 6 weeks and technical meetings when required to resolve issues arising. Early access to the cloud-computing environment with detailed documentation helped to get the nowcast models running first, followed by the SPIDER, then GorgonOps and finally the machine learning models. While it was relatively straightforward to create code to run on an internal computer system, it required a high level of experimentation and expertise to implement this to robustly run in an online operational-like environment. The final move to an autonomous operational system with validated outputs (i.e., at AUL9) is an on-going process.

9. Conclusions

Operational space weather nowcasting and forecasting systems are vital for informing interested stakeholders such as government and infrastructure owners of the potential for hazardous conditions which may affect the integrity of technological systems. The SAGE project was a 4 year project with the aim of developing a series of products to compute the real-time and forecast effects of GICs on critical national infrastructure for GB. The original scientific research models, starting at various levels of maturity, were brought to operational use via a Docker implementation on a cloud-compute system (AWS) for the UK MOSWOC. The models are driven by real-time geomagnetic observatory data or solar wind parameters measured by L1 satellites and accessed via a dedicated API. The products are updated at a 5 min cadence to maintain a balance between processing and data lag times.

Fifteen separate space weather products were created including GIC estimates in the high voltage electricity network, the high pressure gas pipeline network and the railway network of Britain. Within these, four machine learning models were developed to make forecasts of external magnetic field variation at ground level in the UK, substorm and sudden commencement probabilities and the largest rate-of-change of magnetic field within the next 4 hr. Finally, a physics-based MHD model of the near-Earth magnetosphere environment, driven by solar wind measurements, was implemented as an operational real-time system delivering cutting-edge capability in magnetic field forecasting.

Along with the science models and products, three new science-grade variometers were installed across the UK to improve longitudinal coverage of magnetic field measurements. A 4 year campaign to collect magnetotelluric measurements across the island of GB allowed a new 3D conductivity model and magnetic to geoelectric transfer functions to be created. These allow real-time measurements and forecast magnetic field estimates to be convolved to create maps of the geoelectric field across Britain.

Over the 4 years of the project the Application User Levels (AULs) were raised from AUL4 to AUL8; that is, from proven scientific models to pre-operational code running in a representative computing environment. The SAGE program is an excellent example of the R2O pipeline: involving scientific and academic organizations in the implementation of their published models into operational products to inform industrial, commercial and government users of space weather hazard to grounded infrastructure.

Acronyms

API	Application Programming Interface
AUL	Application Usability Level
AWS	Amazon Web Services
GIC	Geomagnetically Induced Currents
HPC	High Performance Computing
MOSWOC	Met Office Space Weather Operations Centre
PSP	Pipe-to-soil Potential
R2O	Research to Operation
SPIDER	Spatial Information from Distributed Exogenous Regression
SAGE	SWIMMR Activities in Ground Effects
SWIMMR	Space Weather Instrumentation, Measurement, Modeling and Risk

Data Availability Statement

Satellite data from ACE and DSCOVR at the L1 Lagrange point were collected from the Space Weather Prediction Center at the US National Oceanic and Atmospheric Administration (NOAA SWPC) website:

https://services.swpc.noaa.gov/products/solar-wind/. Geomagnetic data from the BGS UK observatories are available from the INTERMAGNET website: https://www.intermagnet.org. Magnetotelluric data are published in Hübert et al. (2025). The GB high voltage grid was derived from the Electricity Ten year Statement (Appendix B): https://www.neso.energy/publications/electricity-ten-year-statement-etys. Railway line shapefiles can be sourced at Ordnance Survey: https://osdatahub.os.uk/downloads/open and pipeline locations are available from National Gas: https://www.nationalgas.com/our-businesses/network-route-maps.

Acknowledgments

This study was funded under the NERC SWIMMR N4 programme NE/V002694/1. We would like to thank the following people for their contribution to SAGE: Robert Shore (formerly at BAS), Ellen Clarke (BGS), Eliot Eaton (formerly at BGS), Eleanor Maume (formerly at BGS), Adam Collins (BGS), Guanren Wang (BGS), Helen Smith (BGS), Dave Morgan (BGS), Aideliz Montiel Alvarez (Univ. Edinburgh), Fiona Kirton, (formerly at BGS), Sam Henderson (formerly at BGS), Oli Chambers (BGS), Peter Stevenson (formerly at BGS), Sarah Reay (BGS), Tom Martyn (BGS), Rob Lyon (BGS), Josie Parrianen (BGS), Chris Turbitt (BGS) and Kathryn Whaler (Univ. Edinburgh). We also wish to thank Ian McCrea (RAL Space) for his support as programme chair of SWIMMR. This paper is published with the permission of the Executive Director of the British Geological Survey (UKRI).

References

Beggan, C. D., Richardson, G., Baillie, O., Hübert, J., & Thomson, A. (2021). Geoelectric field measurement, modelling and validation during geomagnetic storms in the UK. *Journal of Space Weather and Space Climate*, *11*, 37. https://doi.org/10.1051/swsc/2021022

Beggan, C. D. (2015). Sensitivity of geomagnetically induced currents to varying auroral electrojet and conductivity models. *Earth Planets and Space*, *67*(1), 1–12. https://doi.org/10.1186/s40623-014-0168-9

Beggan, C. D., Eastwood, J., Forsyth, C., Freeman, M., Heyns, M., Huebert, J., et al. (2024). SAGE: Final report on the AWS nowcast and forecasting system and research advances (Tech. Rep. No. OR/24/013). *British Geological Survey, Open Report*. Retrieved from https://nora.nerc.ac.uk/id/eprint/537731/1/OR24013.pdf

Boteler, D. H. (2013). A new versatile method for modelling geomagnetic induction in pipelines. *Geophysical Journal International*, *193*(1), 98–109. https://doi.org/10.1093/gji/ggs113

Boteler, D. H. (2019). A 21st century view of the March 1989 magnetic storm. *Space Weather*, *17*(10), 1–15. https://doi.org/10.1029/2019SW002278

Butler, E. C., & Keller, J. M. (2021). R2O2R improvements identified by United States space weather forecasters. *Space Weather*, *19*(6), e2021SW002739. https://doi.org/10.1029/2021SW002739

Campanya, J., Gallagher, P. T., Blake, S. P., Gibbs, M., Jackson, D., Beggan, C. D., et al. (2019). Modeling geoelectric fields in Ireland and the UK for space weather applications. *Space Weather*, *17*(2), 216–237. https://doi.org/10.1029/2018SW001999

Campbell, W. H. (1980). Observation of electric currents in the Alaskan oil pipeline resulting from auroral electrojet current sources. *Geophysical Journal of the Royal Astronomical Society*, *61*(2), 437–449. https://doi.org/10.1111/j.1365-246X.1980.tb04325.x

Cannon, P. S. (2013). Extreme space weather — A report published by the UK Royal Academy of Engineering. *Space Weather*, *11*(4), 138–139. https://doi.org/10.1002/swe.20032

Clarke, E., Baillie, O., Reay, S., & Turbitt, C. (2013). A method for the near real-time production of quasi-definitive magnetic observatory data. *Earth Planets and Space*, *65*(11), 1363–1374. https://doi.org/10.5047/eps.2013.10.001

Díaz, J. (2024). Monitoring May 2024 solar and geomagnetic storm using broadband seismometers. *Scientific Reports*, *14*(1), 30066. https://doi.org/10.1038/s41598-024-81079-6

Eggington, J. W. B., Coxon, J. C., Shore, R. M., Desai, R. T., Mejnertsen, L., Chittenden, J. P., & Eastwood, J. P. (2022). Response timescales of the magnetotail current sheet during a geomagnetic storm: Global MHD simulations. *Frontiers in Astronomy and Space Sciences*, *9*, 966164. https://doi.org/10.3389/fspas.2022.966164

Elvidge, S., & Themens, D. R. (2025). The probability of the May 2024 geomagnetic superstorm. *Space Weather*, *23*(1), e2024SW004113. https://doi.org/10.1029/2024SW004113

Eroshenko, E., Belov, A., Boteler, D., Gaidash, S., Lobkov, S., Pirjola, R., & Trichtchenko, L. (2010). Effects of strong geomagnetic storms on northern railways in Russia. *Advances in Space Research*, *46*(9), 1102–1110. https://doi.org/10.1016/j.asr.2010.05.017

Freeman, M. P., Forsyth, C., & Rae, I. J. (2019). The influence of substorms on extreme rates of change of the surface horizontal magnetic field in the United Kingdom. *Space Weather*, *17*(6), 827–844. https://doi.org/10.1029/2018SW002148

Gannon, J. L., Morley, S., Lugaz, N., Liu, H., Carter, B., & Zou, S. (2023). Long-term support is needed for crucial ground-based sensor networks. *Space Weather*, *21*(5), e2023SW003529. https://doi.org/10.1029/2023SW003529

Grandin, M., Bruus, E., Ledvina, V. E., Partamies, N., Barthelemy, M., Martinis, C., et al. (2024). The geomagnetic superstorm of 10 May 2024: Citizen science observations. *EGUsphere*, *2024*, 1–32. https://doi.org/10.5194/egusphere-2024-2174

Halford, A. J., Kellerman, A. C., Garcia-Sage, K., Klenzing, J., Carter, B. A., McGranaghan, R. M., et al. (2019). Application Usability Levels: A framework for tracking project product progress. *J. Space Weather Space Clim.*, *9*, A34. https://doi.org/10.1051/swsc/2019030

Hapgood, M., Angling, M. J., Attrill, G., Bisi, M., Cannon, P. S., Dyer, C., et al. (2021). Development of space weather reasonable worst-case scenarios for the UK national risk assessment. *Space Weather*, *19*(4), e2020SW002593. https://doi.org/10.1029/2020SW002593

Hübert, J., Beggan, C. D., Richardson, G. S., Gomez-Perez, N., Collins, A., & Thomson, A. W. P. (2024). Validating a UK geomagnetically induced current model using differential magnetometer measurements. *Space Weather*, *22*(2), e2023SW003769. https://doi.org/10.1029/2023SW003769

Hübert, J., Beggan, C. D., Richardson, G. S., Martyn, T., & Thomson, A. W. P. (2020). Differential magnetometer measurements of geomagnetically induced currents in a complex high voltage network. *Space Weather*, *18*(4), e2019SW002421. https://doi.org/10.1029/2019SW002421

Hübert, J., Eaton, E., & Beggan, C. D. (2024). *Developing a new ground electric field model for the UK based on long-period magnetotelluric data for the SWIMMR N4*. SAGE framework. Retrieved from http://nora.nerc.ac.uk/id/eprint/537960/BGSOpenReport:OR/24/022

Hübert, J., Eaton, E., Beggan, C. D., Collins, A., & Wang, G. (2025). Long-period magnetotelluric data collected at 44 sites in Scotland, England and Wales [Dataset]. *NERC EDS National Geoscience Data Centre*. https://doi.org/10.5285/14274b67-86a5-4d9d-b3f4-20c6d20228d7

Ingham, M., Divett, T., Rodger, C., & Sigley, M. (2022). Impacts of GIC on the New Zealand gas pipeline network. *Space Weather*, *20*(12), e2022SW003298. https://doi.org/10.1029/2022SW003298

Ingham, M., & Rodger, C. J. (2018). Telluric field variations as drivers of variations in cathodic protection potential on a natural gas pipeline in New Zealand. *Space Weather*, *16*(9), 1396–1409. https://doi.org/10.1029/2018SW001985

Kelly, G. S., Viljanen, A., Beggan, C. D., & Thomson, A. W. P. (2017). Understanding GIC in the UK and French high-voltage transmission systems during severe magnetic storms. *Space Weather*, *15*(1), 99–114. https://doi.org/10.1002/2016SW001469

Lawrence, E., Richardson, G., Reay, S., Thompson, V., Clarke, E., Orr, L., et al. (2025). The geomagnetic and geoelectric response to the May 2024 geomagnetic storm in the United Kingdom. *Frontiers in Astronomy and Space Sciences*, *12*, 1550923. https://doi.org/10.3389/fspas.2025.1550923

Maimaiti, M., Kunduri, B., Ruohoniemi, J. M., Baker, J. B. H., & House, L. L. (2019). A deep learning based approach to forecast the onset of magnetic substorms. *Space Weather*, *17*(11), 1534–1552. https://doi.org/10.1029/2019SW002251

Malone-Leigh, J., Campanyà, J., Gallagher, P. T., Neukirch, M., Hogg, C., & Hodgson, J. (2023). Nowcasting geoelectric fields in Ireland using magnetotelluric transfer functions. *Journal of Space Weather and Space Climate*, *13*, 6. https://doi.org/10.1051/swsc/2023004

Matzka, J., Stolle, C., Yamazaki, Y., Bronkalla, O., & Morschhauser, A. (2021). The geomagnetic Kp index and derived indices of geomagnetic activity. *Space Weather*, *19*(5), e2020SW002641. https://doi.org/10.1029/2020SW002641

Mejnertsen, L., Eastwood, J. P., Hietala, H., Schwartz, S. J., & Chittenden, J. P. (2018). Global MHD simulations of the Earth's bow shock shape and motion under variable solar wind conditions. *Journal of Geophysical Research: Space Physics*, *123*(1), 259–271. https://doi.org/10.1002/2017JA024690

Parsons, A., Biesecker, D., Odstrcil, D., Millward, G., Hill, S., & Pizzo, V. (2011). Wang-Sheeley-Arge–Enlil cone model transitions to operations. *Space Weather*, *9*(3). https://doi.org/10.1029/2011SW000663

Patterson, C. J., Wild, J. A., Beggan, C. D., Richardson, G. S., & Boteler, D. H. (2024). Modelling electrified railway signalling misoperations during extreme space weather events in the UK. *Scientific Reports*, *14*(1), 1583. https://doi.org/10.1038/s41598-024-51390-3

Patterson, C. J., Wild, J. A., & Boteler, D. H. (2023). Modeling "wrong side" failures caused by Geomagnetically Induced Currents in electrified railway signaling systems in the UK. *Space Weather*, *21*(12), e2023SW003625. https://doi.org/10.1029/2023SW003625

Pulkkinen, A., Rastätter, L., Kuznetsova, M., Singer, H., Balch, C., Weimer, D., et al. (2013). Community-wide validation of geospace model ground magnetic field perturbation predictions to support model transition to operations. *Space Weather*, *11*(6), 369–385. https://doi.org/10.1002/swe.20056

Pulkkinen, A., Viljanen, A., Pajunpaa, K., & Pirjola, R. (2001). Recordings and occurrence of geomagnetically induced currents in the Finnish natural gas pipeline network. *Journal of Applied Geophysics*, *48*(4), 219–231. https://doi.org/10.1016/S0926-9851(01)00108-2

Shore, R. M., Freeman, M. P., Coxon, J. C., Thomas, E. G., Gjerloev, J. W., & Olsen, N. (2019). Spatial variation in the responses of the surface external and induced magnetic field to the solar wind. *Journal of Geophysical Research: Space Physics*, *124*(7), 6195–6211. https://doi.org/10.1029/2019JA026543

Smith, A. W., Forsyth, C., Rae, I., Garton, T., Bloch, T., Jackman, C., & Bakrania, M. (2021). Forecasting the probability of large rates of change of the geomagnetic field in the UK: Timescales, horizons, and thresholds. *Space Weather*, *19*(9), e2021SW002788. https://doi.org/10.1029/2021SW002788

Smith, A. W., Forsyth, C., Rae, I., Garton, T., Jackman, C., Bakrania, M., et al. (2022). On the considerations of using near real time data for space weather hazard forecasting. *Space Weather*, *20*(7), e2022SW003098. https://doi.org/10.1029/2022SW003098

Smith, A. W., Freeman, M. P., Rae, I. J., & Forsyth, C. (2019). The influence of sudden commencements on the rate of change of the surface horizontal magnetic field in the United Kingdom. *Space Weather*, *17*(11), 1605–1617. https://doi.org/10.1029/2019SW002281

Smith, A. W., Rae, I., Forsyth, C., Oliveira, D., Freeman, M., & Jackson, D. (2020). Probabilistic forecasts of Storm Sudden Commencements from interplanetary shocks using machine learning. *Space Weather*, *18*(11), e2020SW002603. https://doi.org/10.1029/2020SW002603

Smith, A. W., Forsyth, C., Rae, J., Rodger, C. J., & Freeman, M. P. (2021). The impact of sudden commencements on ground magnetic field variability: Immediate and delayed consequences. *Space Weather*, *19*(7), e2021SW002764. https://doi.org/10.1029/2021SW002764

Thomson, A. W. P., McKay, A. J., Clarke, E., & Reay, S. J. (2005). Surface electric fields and geomagnetically induced currents in the Scottish Power grid during the 30 October 2003 geomagnetic storm. *Space Weather*, *3*(11), S11002. https://doi.org/10.1029/2005SW000156