

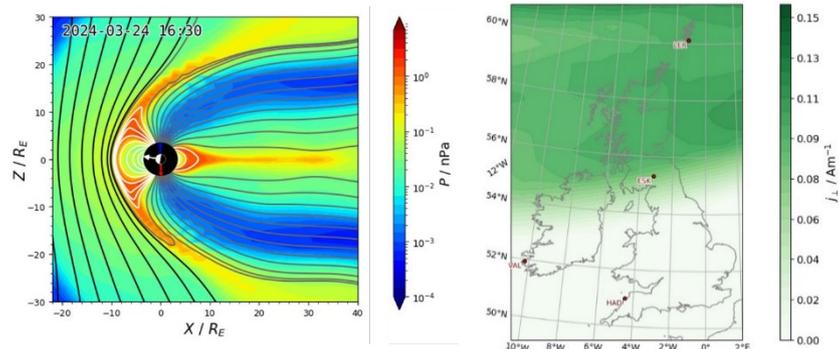


British
Geological
Survey

SAGE: Final report on the AWS nowcast and forecasting system and research advances

Multihazards and Resilience Programme

Open Report OR/24/013



BRITISH GEOLOGICAL SURVEY

MULTIHAZARDS AND RESILIENCE PROGRAMME

OPEN REPORT OR/24/013

Keywords

Report; SWIMMR, SAGE,
Space Weather.

Front cover

Modelled pressure in the magnetosphere and current density over the UK derived from the GorgonOps MHD model. © Imperial College 2024

Bibliographical reference

Beggan, C., et al., 2024.
SAGE: Final report on the AWS nowcast and forecasting system and research advances. British Geological Survey Open Report, OR/24/013. 55pp.

Copyright in materials derived from the British Geological Survey's work is owned by UK Research and Innovation (UKRI) and/or the authority that commissioned the work. You may not copy or adapt this publication without first obtaining permission. Contact the BGS Intellectual Property Rights Section, British Geological Survey, Keyworth, email ipr@bgs.ac.uk. You may quote extracts of a reasonable length without prior permission, provided a full acknowledgement is given of the source of the extract.

SAGE: Final report on the AWS nowcast and forecasting system and research advances

C. D. Beggan, J. Eastwood (Imperial College),
C. Forsyth (MSSL/UCL), M. Freeman (BAS),
M. Heyns (Imperial College), J. Huebert, G.S. Richardson,
A. Smith (MSSL/UCL)

BRITISH GEOLOGICAL SURVEY

Our range of publications is available from BGS shops at Nottingham, London and Cardiff (Welsh publications only). Shop online at shop.bgs.ac.uk.

We publish an annual catalogue of our maps and other publications; this catalogue is available online or from BGS shops.

The British Geological Survey carries out the geological survey of Great Britain and Northern Ireland (the latter as an agency service for the government of Northern Ireland), and of the surrounding continental shelf, as well as basic research projects. It also undertakes programmes of technical assistance in geology in low- to middle-income countries.

The British Geological Survey is a component body of UK Research and Innovation.

British Geological Survey offices

**Nicker Hill, Keyworth,
Nottingham NG12 5GG**

Tel 0115 936 3100

BGS Central Enquiries Desk

Tel 0115 936 3143

email enquiries@bgs.ac.uk

BGS Sales

Tel 0115 936 3241

email sales@bgs.ac.uk

**The Lyell Centre, Research Avenue South,
Edinburgh EH14 4AP**

Tel 0131 667 1000

**Natural History Museum, Cromwell Road,
London SW7 5BD**

Tel 020 7589 4090

Tel 020 7942 5344/45

email bgs-londonstaff@bgs.ac.uk

**Cardiff University, Main Building, Park Place,
Cardiff CF10 3AT**

Tel 029 2167 4280

**Maclean Building, Crowmarsh Gifford,
Wallingford OX10 8BB**

Tel 01491 838800

**Geological Survey of Northern Ireland, Department for
the Economy, Dundonald House, Upper Newtownards
Road, Ballymiscaw, Belfast, BT4 3SB**

Tel 0289 038 8462

www2.bgs.ac.uk/gsni/

**Natural Environment Research Council, Polaris House,
North Star Avenue, Swindon SN2 1EU**

Tel 01793 411500

Fax 01793 411501

www.nerc.ac.uk

**UK Research and Innovation, Polaris House,
Swindon SN2 1FL**

Tel 01793 444000

www.ukri.org

Website www.bgs.ac.uk

Shop online at shop.bgs.ac.uk

Foreword

This report is the published final output on the SWIMMR Activities on Ground Effects (SAGE), a four-year NERC-funded project to implement space weather research into operational products for the Met Office (MO). The SAGE project is a joint enterprise in partnership with British Antarctic Survey, Imperial College London and University College London at Mullard Space Science Laboratory (MSSL).

The report describes the science and technical implementation of predicted Geomagnetically Induced Currents (GICs) models and their effects on the UK high voltage power network, rail network and high-pressure gas pipeline network. The models are coded in Docker containers on the Met Office Amazon Web Services (AWS) platform and are driven by magnetic field inputs both measurements from the BGS observatories and from forecast models using solar wind data at the L1 Lagrange point through a machine learning technique (SPIDER) or a magnetospheric model (GorgonOps). In addition, three other machine-learning-based products provide contextual information on the size, duration, and exceedance rate of an on-going geomagnetic storm.

Though this report focusses on the AWS forecasting and nowcasting system, the SAGE project also undertook significant magnetotelluric (MT) fieldwork to develop a new 3D ground conductivity model specifically for space weather hazard modelling. A separate report details the work and results from the MT fieldwork and modelling.

Acknowledgements

Over four years of the project (April 2020-March 2024), a large number of individuals in BGS, BAS, Imperial and UCL have contributed to the project. This assistance has been gratefully received at all stages of the study. We would particularly like to thank the following people:

Dr Joseph Eggington (Imperial), Dr Rob Shore (BAS), Ellen Clarke (BGS), Eliot Eaton (BGS), Eleanor Maume (BGS), Adam Collins (BGS), Guanren Wang (BGS), Helen Smith (BGS), Dave Morgan (BGS), Aideliz Montiel-Alvarez (Univ. Edinburgh), Dr Fiona Kirton, (BGS), Sam Henderson (BGS), Oliver Chambers (BGS), Peter Stevenson (BGS), Sarah Reay (BGS), Tom Martyn (BGS), Rob Lyon (BGS), Dr Josie Parrienen (BGS) and Chris Turbitt (BGS).

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://www.swpc.noaa.gov/about-space-weather>

With thanks to Edmund Henley (Met Office) for his constructive comments on an initial draft of the report, and to Lauren Orr (BGS) for scientific reviewing of the final report.

Finally, we gratefully acknowledge the work of Dr Alan Thomson (BGS), who led the project as Principal Investigator from 2020-2022.

Contents

Foreword.....	i
Acknowledgements	i
Contents.....	ii
Summary.....	vi
1 Introduction.....	1
1.1 SWIMMR/SAGE.....	1
1.2 N4 Overview.....	1
1.3 N4 Code Overview	3
2 Software and Infrastructure	4
2.1 Amazon Web Services (AWS) Platform.....	4
2.2 Docker.....	4
3 SWIMMR N4 Model.....	8
3.1 Data – volumes and pipelines	8
3.2 Containers.....	21
3.3 Operational Sequencing	25
4 Guide for Developers.....	28
4.1 Repository	28
4.2 Deployment.....	28
4.3 Local Development.....	29
4.4 Future Improvements	30
5 Validation	32
5.1 Products that CANNOT BE Verified	33
5.2 Products that Are not continuously verifiable	33
5.3 Products that CAN BE CONTINUOUSLY Verified	36
5.4 NOTE On Using Real time solar wind data.....	41
References.....	44

FIGURES

Figure 1: Top level data flow and processing steps for the coupled GIC components of the SAGE N4 system. Existing capability was as of June 2020. BGS © UKRI 2024	2
Figure 2: A pictorial representation of the SWIMMR N4 Docker model running on arbitrary infrastructure (e.g. AWS). The round edged coloured boxes represent containers. The sharp-edged grey boxes represent important data outputs from the nearest container. The red cylinders are the interface with the Met Office API. L1 solar wind data, GorgonOps and BGS magnetic observatory data enter the Met Office API from external sources. BGS © UKRI 2024	3
Figure 3: Example of a Dockerfile for the Shock Impact Assessment/M2 model. Each container in the model has its origins in a Docker image and consequently a Dockerfile. BGS © UKRI 2024.....	5
Figure 4: The folders of the swx-swimmr-n4 repository. BGS © UKRI 2024	6
Figure 5: Available datasets and endpoints in the SWIMMR N4 API. BGS © UKRI 2024	9
Figure 6: Example output response from the Met Office API for the solar wind interplanetary magnetic field (IMF) in Geocentric Solar Magnetic (GSM) or Geocentric Solar Ecliptic (GSE) coordinate frame. BGS © UKRI 2024.....	10
Figure 7: Example call to the BGS magnetic observatory data for Lerwick observatory. BGS © UKRI 2024	11
Figure 8: Computation of geoelectric field using thin-sheet model. (a) Magnetic field variation (in nT) over the UK in the magnetic north (Bx) and east (By) is convolved with a thin-sheet model (b) of the UK conductivity model (in milliSiemens). The output is the geoelectric field in the X and (c) Y components in V/km. BGS © UKRI 2024	12
Figure 9: Example of nowcast plot for GIC in the GB high voltage network. (Upper panel) Map of the three network voltage levels (lines) along with area-scaled coloured circles to indicate GIC magnitude and whether an actionable threshold has been crossed. (Lower panels) External field variation in the (North) X and (East) Y component of the magnetic field at the three UK observatories. BGS © UKRI 2024	13
Figure 10: (left) Nowcast Pipe-to-soil potential (in V) in the gas transmission network. (right) Nowcast of the rail index for a large storm. BGS © UKRI 2024	14
Figure 11: (upper panel) Forecast of GIC in the UK HV network using the magnetic field variation from the SPIDER model. (lower panels) Magnetic field measurements (darker lines) and forecast from SPIDER (coloured lines). Dashed vertical line indicates current time. An interactive text box can provide detailed time and magnetic values. BGS © UKRI 2024.....	15
Figure 12: (left) Forecast from SPIDER of the Pipe-to-soil potential (in V) in the gas transmission network. (right) Forecast of the rail index for a large storm. BGS © UKRI 2024	16
Figure 13: Similar to Figure 11, but for a GorgonOps GIC forecast. (upper panel) Forecast of GIC in the UK HV network using the magnetic field variation from GorgonOps. (lower panels) Magnetic field measurements (darker lines) and forecast from GorgonOps (coloured lines). Dashed vertical line indicates current time (note the addition of the previous 5 minutes of forecast data and the baseline normalization). Rail and PSP visualizations also mimic those produced by SPIDER. BGS © UKRI 2024	17
Figure 14: Example of the forecast occurrence of a substorm based on solar wind parameters. Dotted line indicates the forecast period. BAS © UKRI 2024.....	17
Figure 15: Example of the sudden storm occurrence forecast (fourth panel) along with L1 solar wind velocity (upper panel), IMF total field (t) and vertical (z) strength (second panel) and the probability of a shock in the solar wind (third panel). The lower panel shows the Sym H index. © MSSL 2022	18

Figure 16: Example of the exceedance of the rate of change of the magnetic field. (left) As a time series showing L1 solar, predicted rate of change of the magnetic (R) and three separate outputs from various NN models. (right) Visualisation of a snapshot showing concentric circles of probability exceedance. © MSSL 2022	19
Figure 17: GorgonOps estimate of dynamic pressure (in nPascal) within the magnetosphere. © Imperial College 2024	20
Figure 18: GorgonOps estimate of auroral current density over the UK. © Imperial College 2024	20
Figure 19 Flow diagram showing all the different inputs and outputs in the SAGE project. BGS © UKRI 2024	21
Figure 20: Sequence of processing for the nowcast models. BGS © UKRI 2024.....	25
Figure 21: Sequence of processing for the machine learning forecast models. BGS © UKRI 2024.....	26
Figure 22: Sequence of processing for the GorgonOps forecast models. BGS © UKRI 2024...	27
Figure 23: Measured (left axis, blue) and modelled (right axis, black) line GIC for sites Abbey St Bathans (ASB) [East Lothian], Ilderton (ILD) [Northumberland], Knayton (KNA) [North Yorkshire] and Whiteadder (WHI) [East Lothian] for the 4-6 Nov 2018. The zero level for data and model are offset for clarity. Correlation coefficient and polynomial fit coefficients are shown in the bottom left corner of each panel. BGS © UKRI 2021.....	34
Figure 24: ROC curves from a model trained on the full EPT90 (substorm) list from 1996-2014, and curves computed from subsets of the EPT90 test set, for specific waiting times. BAS © UKRI 2022	35
Figure 25: Comparison of the SPIDER model and uncertainty against measured data at the three UK observatories for the 18-Sep-2023 for 19:00-20:00 UT. BGS © UKRI 2024	36
Figure 26: GorgonOps benchmarking relative to the Weimer '05 empirical model given steady state driving, focusing here on ionospheric quantities of interest. © Imperial College 2022	37
Figure 27: GorgonOps benchmarking relative to various empirical models for magnetospheric quantities of interest. © Imperial College 2022	38
Figure 28: Forecast performance for GorgonOps during a G3 geomagnetic storm as measured at Hartland, with the By component of the magnetic field shown here. (Top panel) Measured real-time magnetometer data from the given station. (2nd panel) The full GorgonOps estimated profile, i.e., as if continuously run from L1 data ingested. The solar wind speed which informs the time horizon for a given simulation segment is added to this panel. Both profiles can be thought of as the raw profiles in the past. Note that the measured real-time field is used to estimate a baseline (at t=0) and 5 minutes prior to the forecast time is included in visualisation. (3rd panel) Error associated with each forecast (produced every 5 mins) for the given lead time determined by the solar wind speed. Heat map gives the error between predicted and measured for each subset. Here we can see the structures propagate, but do bear in mind that this shift is naive, and we are aiming to get the last 5 minutes as close as we can (but it is useful to see the profile for the forecast window). (4th panel) RMSE measures for the forecast window, last 5 minutes of the forecast window, and a shifted forecast that assumes the propagation using wind speed is not accurate (see 6th panel). (5th panel) Correlation coefficients for the forecast window and shifted forecast. (6th panel) Assuming the propagation using solar wind speed is not great, the lag estimated using cross correlation is shown here. This is used to shift the forecast in the above panels. © Imperial College 2024	39
Figure 29: GorgonOps ground geomagnetic field estimation considerations. © Imperial College 2024.....	41
Figure 30: Comparison of the plasma moments in the Near-Real-Time (NRT) data with the science quality data. Left column (a, c) shows Advanced Composition Explorer and between 1999 and 2015 (inclusive), right column (b, d) shows Deep Space Climate	

Observatory between 2018 and mid-2019. The plasma moments shown are the plasma density (a, and b) and bulk velocity (c, and d). The diagonal dashed line indicates where the data returned are equivalent, while the orange dashed lines indicate the region where the NRT values are within $\pm 10\%$ of the science data. © MSSL/UCL 2022 42

Figure 31: Fraction of data windows that are continuous, without data gaps for the Advanced Composition Explorer (top: a and b) and Deep Space Climate Observatory (bottom: c, d) Near-Real-Time data. The results for the magnetic field (left: a and c) and plasma data (right: b and d) are shown. The fraction of complete data windows are provided as a function of input window length required to be continuous. Three different interpolation schemes are presented: no interpolation (green circles), interpolation of gaps 5 min or shorter (orange stars) and interpolation of gaps 15 min or shorter (blue crosses). © MSSL/UCL 2022 43

TABLES

Table 1: Summary of products and their verification potential..... 32

Table 2: Evaluation of metrics for 18-Sep-2023 for 19:00-20:00 UT. 37

Table 3: GorgonOps configuration..... 40

Table 4: Comparison of line-of-sight adapted CalcDeltaB (Biot-Savart integration) and Complex Image Method (CIM) methods..... 41

Summary

This report describes the inputs, processing and outputs of the SWIMMR N4 (SAGE) Met Office Amazon Web Services (AWS) Elastic Cloud Compute (EC2) system.

The SAGE system ingests real time magnetic measurements from the British Geological Survey's (BGS) UK geomagnetic observatories and solar wind data measured by satellites at the L1 Lagrange point, provided by the NOAA Space Weather Prediction Centre (SWPC). The data are captured and disseminated by through a Met Office SWIMMR database via a password-protected API.

Four different machine learning models process the L1 solar wind data to produce: (i) a magnetic field forecast for up to one hour ahead of time at the three UK observatories, known as SPIDER; (ii) the probability of a substorm occurring in the next hour (named M1 or the Substorm Forecast); (iii) the probability of a sudden storm commencement in the next hour (denoted M2 or Shock Impact Assessment); and (iv) the probability of the magnetic field rate of change exceeding a series of threshold values at each of the UK observatories in the next hour (called M3 or Extreme Threshold Exceedance Forecast).

A magnetohydrodynamic (MHD) model of the magnetosphere, driven by L1 solar wind data, is run separately on a Met Office HPC system before its data are passed into the SWIMMR API. The code, called GorgonOps, produces magnetic field estimates at the three UK observatories as well as maps of the ionospheric current systems for visualisation purposes. GorgonOps is the operational version of the more general Gorgon global magnetosphere model, and the two names are used interchangeably in this document.

The magnetic field data – either measured for nowcasting or modelled for forecasting – are convolved with a ground conductivity model to produce a set of geoelectric field maps based on the rate of change of the magnetic field which induces a geoelectric field. The snapshot geoelectric field map of Britain is applied to models of (i) the high voltage power grid, (ii) high pressure gas network and (iii) model of the railway system to compute the Geomagnetically Induced Currents (GICs) flowing through them. The model outputs are delivered back to the Met Office API via an automated collection process on the AWS.

The processing is set to run every five minutes, assuming that nowcast and forecast magnetic field data from GorgonOps, and L1 solar wind data are available to the Met Office API (which feeds into SPIDER/Substorm Forecast, Shock Impact Assessment and Extreme Threshold Exceedance Forecast models).

The code to process the data has been set up via docker and docker-compose, allowing modular changes to be made in future. The system is provided as a Github repository which provides full version control and integration into a process that allows the build and deploy to AWS to occur automatically upon triggering a Github commit.

Finally, the report describes the initial validation of the models either within AWS or offline.

1 Introduction

1.1 SWIMMR/SAGE

The SWIMMR (Space Weather Instrumentation, Measurement, Modelling and Risk) project is a £20 million, four-year programme to improve the UK's capabilities for space weather monitoring and prediction. It has an emphasis on space radiation, which can affect aircraft systems, changes in the upper atmosphere affecting communications, and surges in ground currents in power grids and other surface-level systems. These are significant risks to the infrastructures we rely on in daily life and are recorded in the UK's National Risk Register. The programme is funded by UKRI jointly through STFC and NERC.

SWIMMR has a number of high-level objectives including the mitigation of the potential risks of space weather to electric power distribution. A major goal is to produce an updated and improved set of software products and services for forecasting the impact of Geomagnetically Induced Currents (GICs) on power grids and make them available through Met Office Space Weather Operations Centre (MOSWOC). In addition, a second goal is to improve and operationalise the current suite of scientific models for predicting the evolution of the solar wind from the Sun to the L1 Lagrange point.

The SWIMMR activities in ground effects (SAGE) project (termed N4, as it was the fourth NERC-funded item) was proposed to meet these goals, with BGS providing relevant nowcasts for Geomagnetically Induced Currents (GICs) in high voltage power lines, railway lines and pipelines. British Antarctic Survey, Imperial College London and University College London have created a series of codes to provide forecasts of magnetic field variation using machine learning and physics-based models.

The overall set of coupled models (referred to as 'SWIMMR N4' or 'the model') is hosted by the Met Office on their chosen platform, presently an Amazon Web Services (AWS) portal. The GIC data are available at the [Met Office N4 portal](#), and through [the Met Office N4 API](#).

The project ran from June 2020 to March 2024.

1.2 N4 OVERVIEW

Prior to the start of the SAGE project, UK space weather nowcasting and forecasting capability for GICs comprised primarily of scientific-level models developed as research models. BGS have operated a real-time magnetic field data collection at geomagnetic observatories and distribution system over the Internet for over 30 years and had developed a real-time GIC prediction web-accessible portal for the high voltage power network.

The aim of SAGE was to operationalise these models using standardised inputs and to produce standardised outputs. For GIC prediction, the inputs are magnetic field data from the real-time ground observatory network and solar wind data from the ACE or DSCOVR satellites at the L1 Lagrange point around 1 million km upstream of the Earth toward the Sun. The measured magnetic field data from the UK observatories are convolved with a model of the subsurface conductivity to generate a map of the estimated geoelectric field experienced on the surface within the past few minutes.

Using the geoelectric field maps, we place models of the high voltage power network, rail network or high-pressure gas pipeline network. This type of infrastructure is electrically connected to the ground either at single points (like electrical substations) or continuously in the case of the pipeline network. In common, they offer a low resistance path for the equalisation of current across long 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 magnetic flux escape from the core and increased reactive power demand causing voltage harmonics and grid instability. Electrified rail lines can experience signalling faults, which depending on the system configuration can give rise to 'right' or 'wrong side' failures,

while pipelines experience excess pipe-to-soil potential which can cause long term damage from increased corrosion rates.

Figure 1 shows a top-level view of the data inputs, outputs and intermediate processing steps.

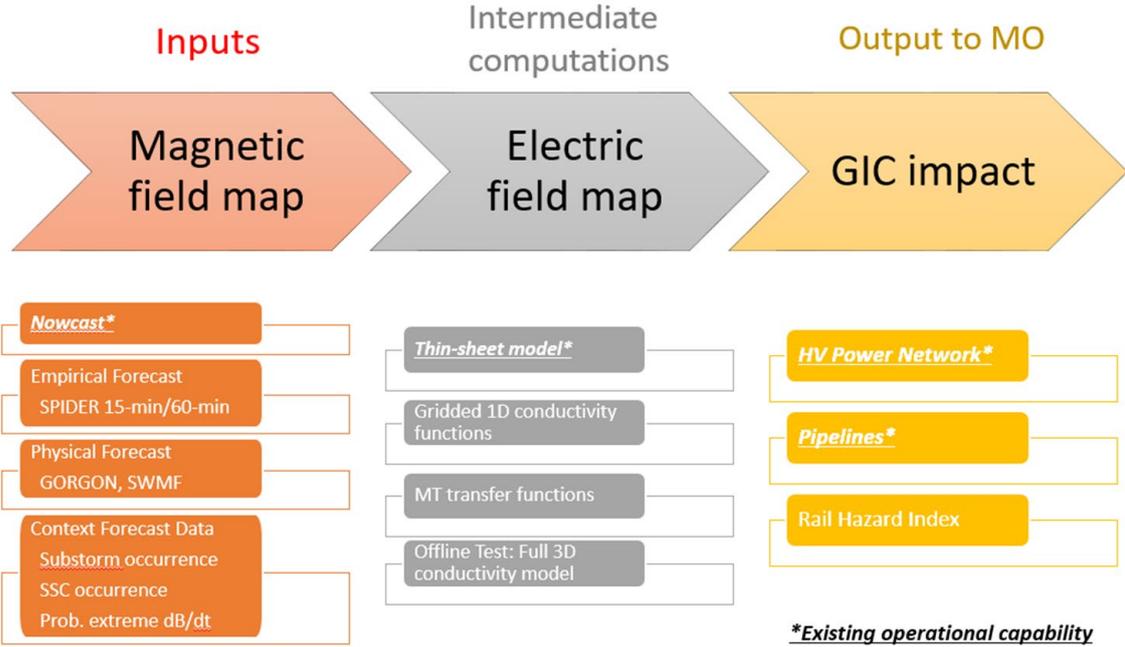


Figure 1: Top level data flow and processing steps for the coupled GIC components of the SAGE N4 system. Existing capability was as of June 2020. BGS © UKRI 2024

In addition to the real time magnetic field measurements, two forecasting models were implemented. The SPIDER and GorgonOps models use machine learning and physics-based modelling, respectively, to determine the variation of the magnetic field on the ground in the UK. These can be combined with the ground observatory data to forecast what the expected geoelectric field and GIC will be up to one hour ahead of time, though less during a severe space weather event with elevated solar wind speeds. As a fallback, the Space Weather Modelling Framework (SWMF) code was also considered by Met Office (noted in Figure 1).

There are also three machine learning-based contextual forecast models developed which provide a longer term (hours to days) estimate of how an on-going geomagnetic storm might behave. The first is a substorm occurrence (labelled Substorm Forecast) model predicts the next occurrence of the Dungey cycle substorm initiation phase. Substorms are the indirectly driven portion of geomagnetic storms, which occur when magnetic reconnection in the tail drives energy back into the ionosphere, energising the auroral oval. The Substorm Forecast model uses solar wind parameters to correlate the loading of the magnetotail but does not predict substorms “from out of the blue”.

The sudden storm commencement (SSC) occurrence model (known as Shock Impact Assessment) estimates if a jump in the solar wind speed and density has occurred and how long a geomagnetic storm may continue for, depending on the size of the SSC.

Finally, the probability of extreme rates of change of the magnetic field (dB/dt) occurring at the three UK observatories is provided by the Extreme Threshold Exceedance Forecast model. It computes a series of estimates of the chance of a predetermined threshold of dB/dt being exceeded at each location.

Further detailed descriptions on each model input and output are given in Sections 2 and 3.

1.3 N4 CODE OVERVIEW

The N4 code product resides in the SWIMMR N4 GitHub private repository – [swx-swimmr-n4](#). The repository (often called a repo) contains all the code and data files required to produce a series of Docker images that contain the instructions and software required to acquire and process the input space weather data. In the next section, the flow of data and information through the AWS system is described. The data inputs, scientific processing and resulting outputs are also discussed. Figure 2 shows an overview of how the Docker passes data between the various parts and process the inputs.

The repository also contains instructions to create runnable instances of the code in images, also known as containers. The resulting containers work in tandem to produce the now and forecast GIC data streams. Note that more than one container can be launched from the same image. More information on each container can be found in Section on Containers.

The Met Office offers a comprehensive set of documentation for a general SWIMMR model. This document complements the existing documentation by going into depth about the SWIMMR N4 models specifically and explains the underlying concepts in more detail.

The first part of the report is written assuming some familiarity with technical concepts git, Linux, and scientific concepts involved in magnetic field measurement, machine learning prediction, geoelectric field and GIC computation.

The second part looks at model validation and the new scientific outputs from the project.

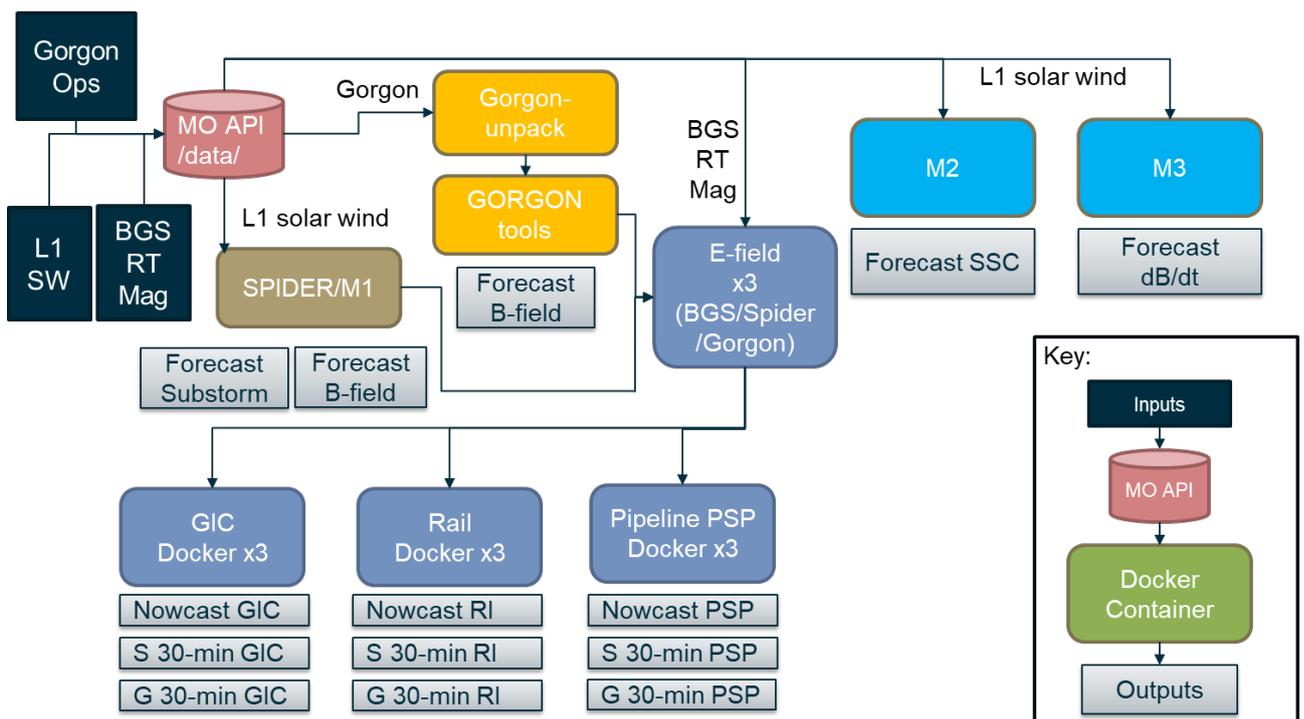


Figure 2: A pictorial representation of the SWIMMR N4 Docker model running on arbitrary infrastructure (e.g. AWS). The round edged coloured boxes represent containers. The sharp-edged grey boxes represent important data outputs from the nearest container. The red cylinders are the interface with the Met Office API. L1 solar wind data, GorgonOps and BGS magnetic observatory data enter the Met Office API from external sources. BGS © UKRI 2024

2 Software and Infrastructure

2.1 AMAZON WEB SERVICES (AWS) PLATFORM

The N4 model has been developed, coded and deployed on an [Amazon Elastic Compute Cloud Server](#) (Amazon EC2). The EC2 server is part of a wider suite of software/hardware within the [Amazon Web Services](#) (AWS) ecosystem, which includes public and private storage options (known as [Buckets](#)) and the [Elastic Container Registry](#) (ECR) which allows scaling of applications and compute requirements.

The SWIMMR N4 repository is located: <https://github.com/MetOffice/swx-swimmr-n4>

The default branch is *develop* which is where code updates are pushed to by the N4 partners.

In December 2023, SWIMMR N4 was running on Amazon Linux 2 with:

- CPE_NAME="cpe:2.3:o:amazon:amazon_linux:2"
- Linux kernel release 4.14.290-217.505.amzn2.x86_64.
- The build environment is [AWS Codebuild](#) using the `aws/codebuild/standard:6.0` image.
- Docker version 20.10.17.
- Docker compose version 2.6.0.

A specific SWIMMR project (S3) run by the Met Office has developed the AWS backend infrastructure. The N4 project also collaborated with S3 on new features that we required such as the implementation of [Docker Secrets](#) (which allows passwords to be exchanged that BGS and partner organisations wish to keep hidden) as well as expanding spaces in the container registry¹.

Contact with the SWIMMR support team is started using the SWIMMR@metoffice.gov.uk e-mail address.

2.2 DOCKER

2.2.1 Introduction

The N4 application is written in the form for a containerised environment. This is to ensure modularity and secure execution of the code. There is a relatively small amount of software required to get started, with most the science and data processing happening inside the containers (view the section Containers). The only required prerequisite piece of software to run the model is [Docker](#), and the N4 model can be run on a variety of infrastructure as long as it supports the Docker client and the Docker daemon. Currently, the project is working with a Docker client and Docker server version of **20.10.17**. The Met Office provides the [Packaging a Model documentation](#) section which includes its own Docker explanations.

2.2.2 Docker Images

The foundational building block of Docker is the Docker image, which can, for our purposes, be considered as a virtual instance of a Linux machine. Images are built using the `docker build` command and images are effectively compiled instructions for creating containers. Note that Docker images do not operate on their own. Images are built from a combination of other images and a controlling [Dockerfile](#) script.

As an example, the contents of one of the Dockerfiles used to create an image and then create a model container are shown in Figure 3. In this case, a model code (M2_ModelRun.py) is built to run under Python3.7 and then executed (using CMD) when the image is built.

¹ The spaces are for Docker image names, not Docker containers.

 ochambers Set up crons and output file structure

Code Blame 22 lines (15 loc) · 527 Bytes

```
1 FROM public.ecr.aws/docker/library/python:3.7
2
3 # Set up other shared environment variables
4 ENV IN_DOCKER=1
5
6 COPY . /root
7 WORKDIR /root
8
9 COPY m2cron /etc/cron.d/
10
11 # Installing system dependencies
12 RUN apt update && apt install -y cron && \
13     apt install --no-install-recommends -y build-essential gcc && \
14     apt clean && rm -rf /var/lib/apt/lists/*
15
16 # Installing python dependencies
17 RUN pip3 install --no-cache-dir -r /root/requirements.txt
18
19 # RUN chmod -x /model
20
21 ENTRYPOINT ["bash", "/root/m2entry.sh"]
22 CMD [ "cron", "-f" ]
```

Figure 3: Example of a Dockerfile for the Shock Impact Assessment/M2 model. Each container in the model has its origins in a Docker image and consequently a Dockerfile. BGS © UKRI 2024

2.2.3 Containers

Docker containers are created from images, and containers perform the required data ingestion, processing and generation of outputs. In the context of the model, containers generally consume and process space-weather data, creating derived outputs. Containers are created using the *docker run* command. Generally, a container has its own storage (virtually isolated from the host machine), and effectively its own operating system and software. Containers can be created from the same image to perform different functions by passing arguments to the *docker run* command or by changing accessible environment variables.

When starting SWIMMR N4 these containers are launched with different functionality by manipulating the environment variables at launch. A container's purpose is generally to perform the command(s) set out in the parent images Dockerfile **CMD** or **ENTRYPOINT** instructions.

For SWIMMR N4, Figure 4 shows how the Docker-Compose containers are set up. As an example, This follows on from Figure 2 which shows how the three Rail Docker containers process 1) geoelectric field data from the BGS nowcast magnetic field values, 2) forecasts from the SPIDER container and 3) forecasts from the Gorgon container.

📁 n4-model-bas-spider
📁 n4-model-bgs-base
📁 n4-model-bgs-efield
📁 n4-model-bgs-gic
📁 n4-model-bgs-pipeline
📁 n4-model-bgs-rail
📁 n4-model-icl-gorgon-tools
📁 n4-model-icl-monsoon-unpack
📁 n4-model-template
📁 n4-model-ucl-m2
📁 n4-model-ucl-m3
📁 n4-utility-bgs-general

Figure 4: The folders of the swx-swimmr-n4 repository. BGS © UKRI 2024

2.2.4 Volumes

By design, containers run in a virtual isolation and cannot use or access the storage of their sibling containers. However, in SWIMMR-N4 it is required that the outputs from some of the containers be used as the inputs to other containers. To facilitate this data sharing, the Met Office recommend the use of [Docker Volumes](#). Volumes are a Docker object that can be mounted at any directory in any container running on the Docker server.

For example, Figure 2 shows that the GIC container needs the output of the E-field container. The E-field (geoelectric field) container already writes its outputs to `/root/efield`. To facilitate the data transfer, the E-field container mounts the volume `efield_calcs` at `/root/efield`. Then, the GIC container mounts the volume `efield_calcs` at `/root/efield`².

Thereafter the GIC container can access all the outputs of the E-field container by looking in `/root/efield`, and writes its outputs to a volume mounted at `/root/gic_results`. The contents of `gic_results` (and similar volumes for the *spider*, *psp*, *rail* and *gorgon* outputs) is then regularly copied to a volume `/root/copy_out/` which is mapped to `/output` where output files are picked up for the SWIMMR portal. They are copied to `copy_out`, rather than mapping directly to `/output` as `/output` is regularly cleared out which can cause issues with files that need to persist long

² In SWIMMR N4, path used for the output mount of the E-field container is the same as the path used for the input mount of the GIC container. Therefore, the GIC container looks for its inputs at `/root/efield`. There is no reason the paths need to be the same.

enough for other code to use them. SWIMMR N4 creates and coordinates Docker volumes by using Docker Compose.

2.2.5 Docker Compose

[Docker Compose](#) is yet another technology within the Docker ecosystem. Unlike the Docker client and server, it is arguably not a prerequisite to running SWIMMR N4, because anything that can be done with the *docker compose* command can be done with other Docker commands. However, for our purposes *docker compose* is essential to the start-up of the model. SWIMMR N4 was successfully running using a Docker compose version of **2.6.0**.

Docker compose is a command line utility that orchestrates the start-up and operation of other Docker components such as containers, volumes, and secrets and optionally the building of images. Its primary purpose is to make complex setups easier.

The fundamental *docker compose* command is *docker compose up*. By default, *docker compose up* carries out the instructions defined in a [docker-compose.yml](#) file located in the same directory as the command is run. At the time of writing, starting SWIMMR N4 on the Met Office AWS infrastructure requires running the command *sudo trigger_model*. The primary operation that *trigger_model* completes is to run *docker compose up*.

2.2.6 Secrets

To ensure that passwords used to access BGS python libraries on a BGS webserver are kept safe within the GitHub repository and can be injected (semi-)securely at build time, the model ingests a [Docker Secret](#) file. A Docker secret takes the form of an encrypted file on the host machine of the Docker server. Docker secrets are plain text files that contain passwords or keys, or scripts that carry out commands or set environment variables.

The BGS Docker secret is a bash script that sets environment variables that are referred to in various Dockerfiles within the model components. It is stored separately by the Met Office elsewhere in the AWS ecosystem (a solution for local development is noted in Local Development). Configuring which images are built with the Secret is defined in the *docker-compose* file.

3 SWIMMR N4 Model

In this section, the flow of data and information through the AWS system is described. The data inputs, scientific processing and resulting outputs are also discussed.

3.1 DATA – VOLUMES AND PIPELINES

3.1.1 Met Office API Overview

The Met Office provide dedicated API access to each of the SWIMMR projects that allows specific types of magnetic and solar wind data to enter into the N4 model. API access is controlled through calls to a password-key protected set of web links (often termed endpoints). These can be programmatically accessed using standard web protocols or Python 3 packages such as *requests*.

Initial documentation was provided at the start of the project that can be loaded into <https://editor.swagger.io/> to check and visually inspect endpoints. The endpoint for inputs is entered through the following general link:

https://gateway.api-management.metoffice.cloud/swx_swimmr_n4/1.0/v1/data/<path-to-input-data>

where the endpoint link is followed by text (<path-to-input-data>) linking to the specific data product. These include solar wind velocity, density, Interplanetary Magnetic Field magnitude and particle temperature at the L1 Lagrange point. These values are ingested by Met Office from the OMNI website as measured by the ACE or DSCOVR satellite missions.

Real-time minute-mean cadence magnetic field values from the three UK observatories operated by BGS are provided to Met Office via a password protected webpage who then make the BGS observatory data available through the API.

The outputs from SWIMMR N4 are made available at a similar endpoint:

https://gateway.api-management.metoffice.cloud/swx_swimmr_n4/1.0/v1/output/<path-to-output-file>

The outputs are also available at a front-end for the API, the SWIMMR N4 [portal](#). More documentation on the API can be found in the [Data Retrieval documentation](#).

3.1.2 Data Inputs

The Met Office API is an access point for a MongoDB database. It holds long records of ingested geophysical data.

3.1.2.1 L1 SPACE WEATHER DATA

The Met Office ingest solar wind data from the [NOAA Space Weather Prediction Centre \(SWPC\)](#) real-time database operated by NOAA. The data are collected by a number of spacecraft but primarily ACE and DSCOVR. Note there are occasional gaps in this dataset when the L1 satellites are not able to downlink data, for example, leading to a delay, or if instruments are offline for maintenance.

3.1.2.2 BGS MAGNETIC DATA

The BGS provide real-time data 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. The data are provided every minute and ingested from the dedicated BGS webpage into the Met Office API. Occasional repeated measurements occur (~1%) which relate to time-lags between completion of the processing of the minute mean values on the BGS server and update of the webpage. The Met Office API notes the magnetic values at each timestamp and their ingestion time. The latest ingested value is recommended, though as the

difference between values is usually only 0.1 nT that is within the uncertainty of the measurement in any case.

3.1.3 Available Met Office API data

The SWIMMR SAGE (N4) API has the following endpoints shown in Figure 5. These provide solar wind parameters and ground magnetic field values for the latest measured values or for a range of times in the past.

SWIMMR SAGE API 0.1.0 OAS 3.0

default ^

GET	/v1/data/rtsw_magnetometer	Returns magnetometer datapoints between two times.	∨
GET	/v1/data/rtsw_magnetometer/latest	Returns latest magnetometer datapoint.	∨
GET	/v1/data/rtsw_wind	Returns solar wind datapoint between two times.	∨
GET	/v1/data/rtsw_wind/latest	Returns latest solar wind datapoint.	∨
GET	/v1/data/rtsw_ephemerides	Returns ephemerides datapoints between two times.	∨
GET	/v1/data/rtsw_ephemerides/latest	Returns latest ephemerides datapoint.	∨
GET	/v1/data/external_rolling_magnetometer	Returns external only, rolling magnetometer data between two times	∨
GET	/v1/data/external_rolling_magnetometer/latest	Returns latest external only, magnetometer value.	∨
GET	/v1/data/uk_rolling_magnetometer	Returns 24 hour magnetic field measurements between two times.	∨
GET	/v1/data/uk_rolling_magnetometer/latest	Returns latest 24 hour magnetic field measurements of a given type.	∨
GET	/v1/output/list	Returns a list of objects from an S3 bucket.	∨
GET	/v1/output/list_directories	Returns a list of directories only from an S3 bucket.	∨
GET	/v1/output	Returns a presigned URL link for a given object.	∨

Figure 5: Available datasets and endpoints in the SWIMMR N4 API. BGS © UKRI 2024

3.1.3.1 L1 SOLAR WIND

Data calls to the Met Office API are programmed in Python and typically returned as JSON structures from which the time-stamped information is extracted. For the solar wind, the outputs are separable into a Python dictionary.

For example, a call to the real-time solar wind (`rtsw_wind/latest`) for a particular timestamp returns the output shown in Figure 6.

```
{
  "active": true,
  "source": "DSCOVR",
  "range": 0,
  "scale": 4,
  "sensitivity": 0.002,
  "manual_mode": false,
  "sample_size": 3000,
  "bt": 4.54,
  "bx_gse": -0.19,
  "by_gse": 4.22,
  "bz_gse": 1.62,
  "theta_gse": 20.94,
  "phi_gse": 92.62,
  "bx_gsm": -0.19,
  "by_gsm": 4.52,
  "bz_gsm": 0.2,
  "theta_gsm": 2.56,
  "phi_gsm": 92.45,
  "max_telemetry_flag": 0,
  "max_data_flag": 0,
  "overall_quality": 0,
  "ingestion_timestamp": "2023-11-27T12:59:50",
  "timestamp": "2023-11-27T12:55:00"
}
```

Figure 6: Example output response from the Met Office API for the solar wind interplanetary magnetic field (IMF) in Geocentric Solar Magnetic (GSM) or Geocentric Solar Ecliptic (GSE) coordinate frame. BGS © UKRI 2024

The call returns magnetic field values (B) and the relative angle of the magnetic field orientation along with time and spacecraft information. In Figure 6, `Bz_gsm` is the magnitude of the vertical component of the interplanetary magnetic field (Bz) in the geocentric solar magnetic (GSM) coordinate frame. Other solar wind parameters are plasma density (n_e), temperature ($^{\circ}\text{K}$), measurement satellite source (DSCOVR/ACE) and ephemerides information (i.e. the location in GSM coordinates).

These values of the solar wind IMF and plasma are used as input to the SPIDER, Substorm Forecast, Shock Impact Assessment, Extreme Threshold Exceedance Forecast and GorgonOps code.

3.1.3.2 BGS B-FIELD

The external only or full field (rolling) values of the three UK observatories can be extracted from the Met Office API. A default call to the `external_only` rolling value returns:

```
https://gateway.api-management.metoffice.cloud/swx\_swimr\_n4/1.0/v1/data/external\_rolling\_magnetometer/latest
```

The latest values come from Lerwick by default (Figure 7) but an additional parameter allows for Hartland or Eskdalemuir to be polled.

These data along with Hartland and Eskdalemuir are used to generate real time geoelectric field maps for the UK.

```

{
  "Bx": -22.2,
  "By": -5.4,
  "Bz": 13.2,
  "timestamp": "2023-11-27T13:02:00",
  "ingestion_timestamp": "2023-11-27T13:04:06",
  "type": "LERWICK"
}

```

Figure 7: Example call to the BGS magnetic observatory data for Lerwick observatory. BGS © UKRI 2024

3.1.3.3 MONSOON INPUT

To facilitate GorgonOps model output transfer from the Met Office Monsoon2 HPC environment, output is uploaded to the SWIMMR N4 S3 bucket in a compressed zipped (.zip) binary format. The *monsoon-unpack* docker container is then able to download the binary file, unpacking and processing the data on AWS and making the output available to other containers.

3.1.4 Container Throughputs

3.1.4.1 ELECTRIC FIELD

The [electric field docker](#) reads magnetic field time-series from either the real-time BGS external field measurements or the SPIDER or GorgonOps forecasts of the magnetic field and convolves the values with a conductivity model.

The original (2022/23) conductivity model is based on the ‘thin-sheet’ 2.5D model of the UK conductivity. The conductivity model is included in the Docker image and is copied from the Git repository when the Docker image is built. The thin-sheet model is computationally expensive so is written in FORTRAN which is compiled on the docker image build. A hook using the Python *f2py* module allows the magnetic field rate of change to be passed into the FORTRAN code.

The magnetic field inputs are created by loading in the variation of the magnetic data from the three observatories using the API. Data from the previous five minutes are loaded, choosing the minute with the large variation at any observatory. The three observatory values are used to create a simple linear extrapolation of the values between each location for both B_x and B_y .

The magnetic values from the observatories are extrapolated across the UK and Ireland to create a 10 x 10 km grid map over the British Isles covering 49°-62°N and -12°-+2°E which is around 1800 x 1300 km. The snapshot map of magnetic field variation is passed to be FORTRAN code which computes an anomalous geoelectric field map of the same size. Figure 8 illustrates the steps of magnetic field convolved with a conductivity model to produce a geoelectric field map. Note (a) and (c) are examples from different time periods and are shown only to illustrate what heightened geomagnetic activity looks like.

The electric field container generates a snapshot map of the *geoelectric* field across the UK mainland at a 10 x 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 Gorgon. The map (either nowcast or forecast) is passed internally within the Docker compose environment to the GIC modules for the high-voltage power grid, railway index and high-pressure pipeline.

The electric field container does not save any data or maps to the /outputs folder for further ingestion back into the API.

Note, future development will replace the thin-sheet model with one based on maps created using around 75 magnetotelluric transfers functions measured in the SWIMMR-funded campaign of 2021-2024.

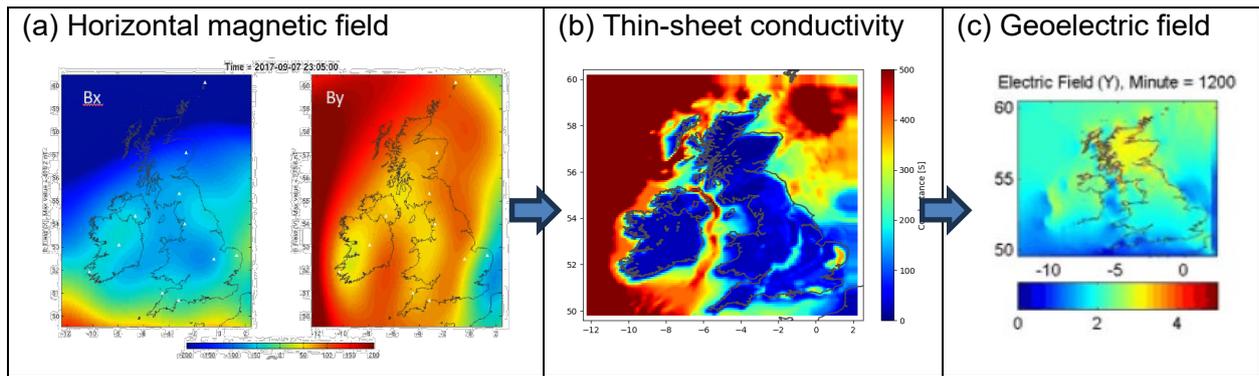


Figure 8: Computation of geoelectric field using thin-sheet model. (a) Magnetic field variation (in nT) over the UK in the magnetic north (Bx) and east (By) is convolved with a thin-sheet model (b) of the UK conductivity model (in milliSiemens). The output is the geoelectric field in the X and (c) Y components in V/km. BGS © UKRI 2024

The electric field container generates a snapshot map of the *geoelectric* field across the UK mainland at a 10 x 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 GorgonOpsOutputs

There are three primary outputs from the GIC containers (HV GIC, Pipeline and Rail) which are the nowcast from the real time measurements of the ground magnetic field and computed geoelectric field, the forecast from the SPIDER machine learning model and the forecast from the GorgonOps physics-based model of the magnetosphere.

The outputs consist of (i) computed GIC in the high voltage transformers on the 400 kV, 275 kV and 132 kV (in Scotland) National Grid power transmission system, (ii) the Pipe-to-Soil Potential (PSP) in the high-pressure gas pipeline and (iii) a rail index indicating where issues in rail signalling might occur in the GB network.

This sub-section explains the outputs as forecasts or nowcasts while Section 3.2 explains the outputs from each individual docker.

3.1.4.2 NOWCASTS

The *estimation of GIC* in the high voltage grid consists of a file of GIC in Ampere 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 an unearthed part of the network (in which case GIC is zero).

The location and magnitude of GICs in the power grid are shown in Figure 9 which provides an example of the nowcast visualisation. The map shows the three network voltage levels (as coloured lines) along with area-scaled coloured circles to indicate GIC magnitude at each substation (in ampere, A). The colours (blue, green, orange and red) indicate whether an actionable threshold has been crossed. For context, the two lower plots show the external field variation in the (North) X and (East) Y component of the magnetic field at the three UK observatories. The table in the upper right shows the largest modelled GIC in the past five minutes and the short-code location, e.g. INDQ is Indian Queens in Cornwall.

As a rule of thumb, GIC larger than 25 A becomes hazardous to certain types of older transformers. A GIC larger than 50 A is likely to begin to cause problems with modern transformers. GIC values larger than 100 A for sustained periods are considered to be very hazardous regardless of transformer type or age. However, what is provided is substation GIC which would be divided between the number of transformers on the given site. This usually between 1 and 8 depending on the substation.

For magnetic field variations in X and Y, at the UK observatories, over the course of 1 hour, a rule of thumb is that variations larger than 50-100 nT are a minor storm, while variations over 250-500 nT are a large storm. Variations over 1000 nT are considered an extreme storm.

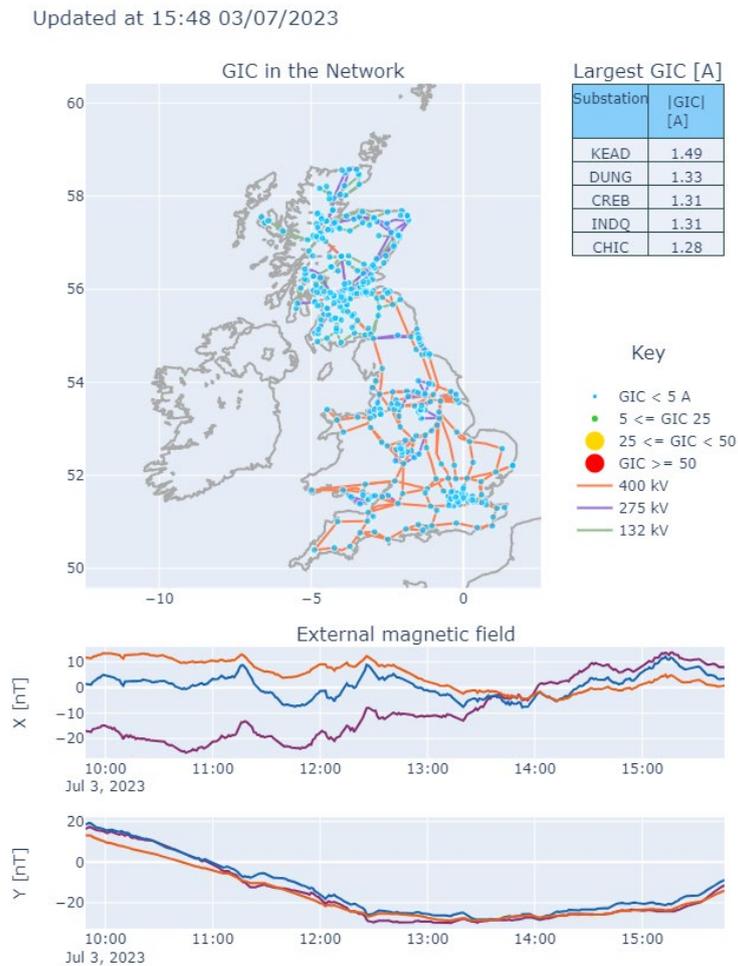


Figure 9: Example of nowcast plot for GIC in the GB high voltage network. (Upper panel) Map of the three network voltage levels (lines) along with area-scaled coloured circles to indicate GIC magnitude and whether an actionable threshold has been crossed. (Lower panels) External field variation in the (North) X and (East) Y component of the magnetic field at the three UK observatories. BGS © UKRI 2024

Pipe-to-soil potential (PSP) is calculated at thousands of points along the GB high-pressure network. The output is a file of position and PSP in volts.

For the visualisation of pipe-to-soil potential (PSP) output, the map in Figure 10 indicating the polarity and magnitude of the parameter in volts (V) as colour intensity. As with the high voltage grid, the PSP estimates are based on the largest variation of the magnetic field (and hence geoelectric field) in the past five minutes.

As a rule of thumb, values outside the range of -0.85 V to -1.2 V are considered to be outside the normal operation of the pipeline.

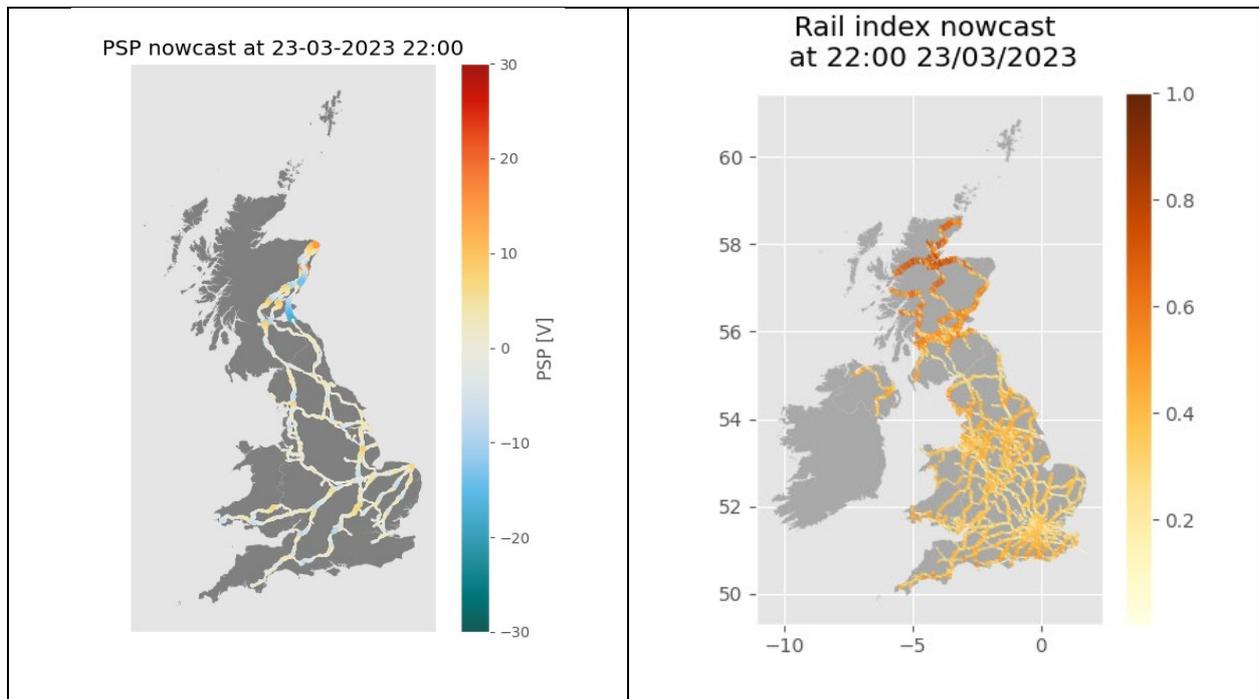


Figure 10: (left) Nowcast Pipe-to-soil potential (in V) in the gas transmission network. (right) Nowcast of the rail index for a large storm. BGS © UKRI 2024

The **rail index** is generated by computing the geoelectric field incident to the direction of the rail track in the 10 x 10 km grid cell. This value is then scaled to the magnitude of the modelled 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 csv file with around 82,000 line sections.

Values above 0.8 indicate the risk of misbehaviour of rail side equipment, though further research is needed as this depends on a large number of factors such as line type, equipment type, age and electrification. Figure 10 (right) illustrates a snapshot map showing the rail index during a large storm.

3.1.4.3 FORECASTS

The forecasts of future magnetic field variation (typically for the next hour) are created from the SPIDER and GorgonOps codes. There are three additional contextual forecasts named Substorm Forecast, Shock Impact Assessment and Extreme Threshold Exceedance Forecast which are machine learning models to predict (1) the probability of a magnetic substorm occurring, (2) the probability of a storm sudden commencement occurring and (3), the probability of the change of the magnetic field exceeding a set of threshold levels at each of the UK observatories in the next few hours. We describe the forecasts in more detail next.

Magnetic field forecast from SPIDER

The **SPIDER model** forecasts between 30 minutes to 60 minutes ahead of time depending on the velocity of the solar wind at the L1 Lagrange point. The solar wind data are collected from the Met Office API and the machine learning model forecasts the external magnetic field variation at the three UK observatories. The forecast of the magnetic field is read by the geoelectric field Docker and a forecast geoelectric field is created. From the forecast geoelectric field, the GIC in the HV GB power grid can be computed (Figure 11, upper panel). The forecast can be compared against or visualised with the measured ground magnetic field up to that point in time (Figure 11, lower panels).

Forecast for 22:00 23/03/2023

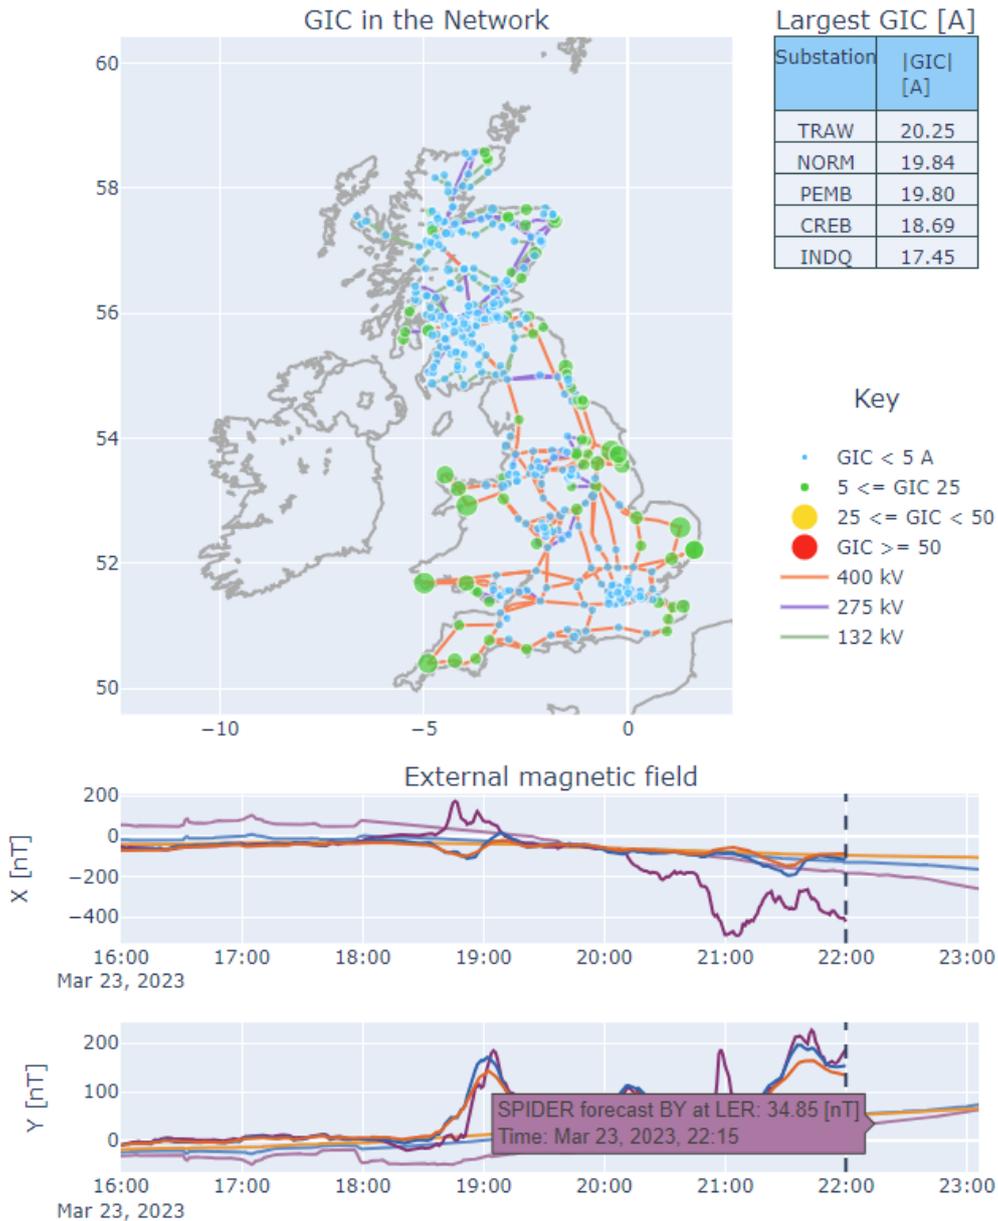


Figure 11: (upper panel) Forecast of GIC in the UK HV network using the magnetic field variation from the SPIDER model. (lower panels) Magnetic field measurements (darker lines) and forecast from SPIDER (coloured lines). Dashed vertical line indicates current time. An interactive text box can provide detailed time and magnetic values. BGS © UKRI 2024

Similarly, the PSP and Rail Index can be forecast up to one hour ahead of time (Figure 12). The output from the GorgonOps can be treated in the same way to predict GIC in the power, pipeline and rail networks.

Magnetic field forecast from GorgonOps

GorgonOps is able to run continuously in real-time in the Monsoon2 HPC environment, ingesting real-time solar wind data from the SWIMMR API. While the model is running, outputs are packed and transferred to the SWIMMR API every 5 minutes. On AWS the *monsoon-unpack* container is run every 5 minutes, unpacking all the relevant files for a given forecast horizon (determined by the current solar wind speed). The resulting profile is then shifted relative to the forecast horizon, with measured geomagnetic field data from the various UK

observatories used to determine the baseline. The resulting ground geomagnetic field datafile for each station includes the forecast for the previous 5 minutes and padded up to the following hour. Simulation and forecast time are included, where the latter is the shifted simulation output relative to the solar wind speed at the current time instance. Where forecast data from Monsoon is available, the ground geomagnetic field estimate along with a status flag is used. In the case of missing data, a 99999.99 placeholder is used.

The output is placed in: /gorgon/yyyy/mm/dd/HH/MM/dB_{ESK,LER,HAD}.csv

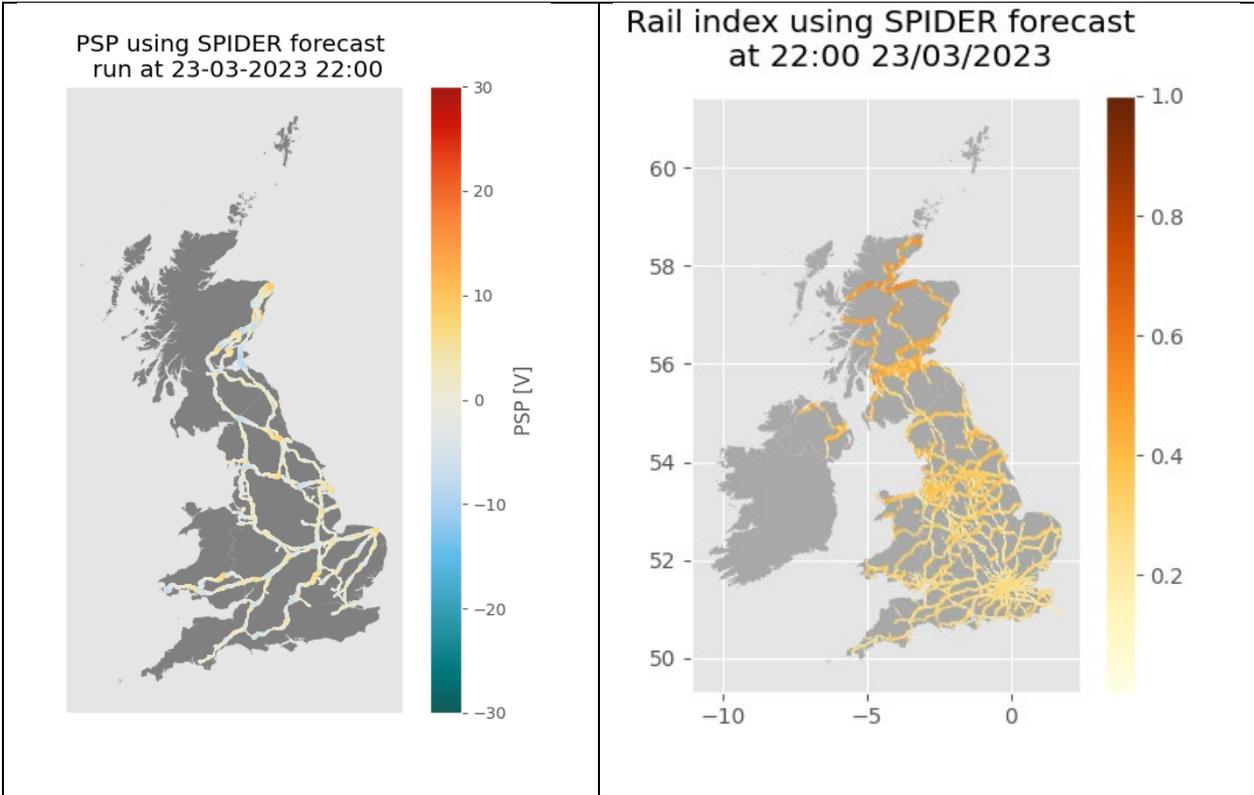


Figure 12: (left) Forecast from SPIDER of the Pipe-to-soil potential (in V) in the gas transmission network. (right) Forecast of the rail index for a large storm. BGS © UKRI 2024

Forecast for 14:23 02/05/2024 (UT)

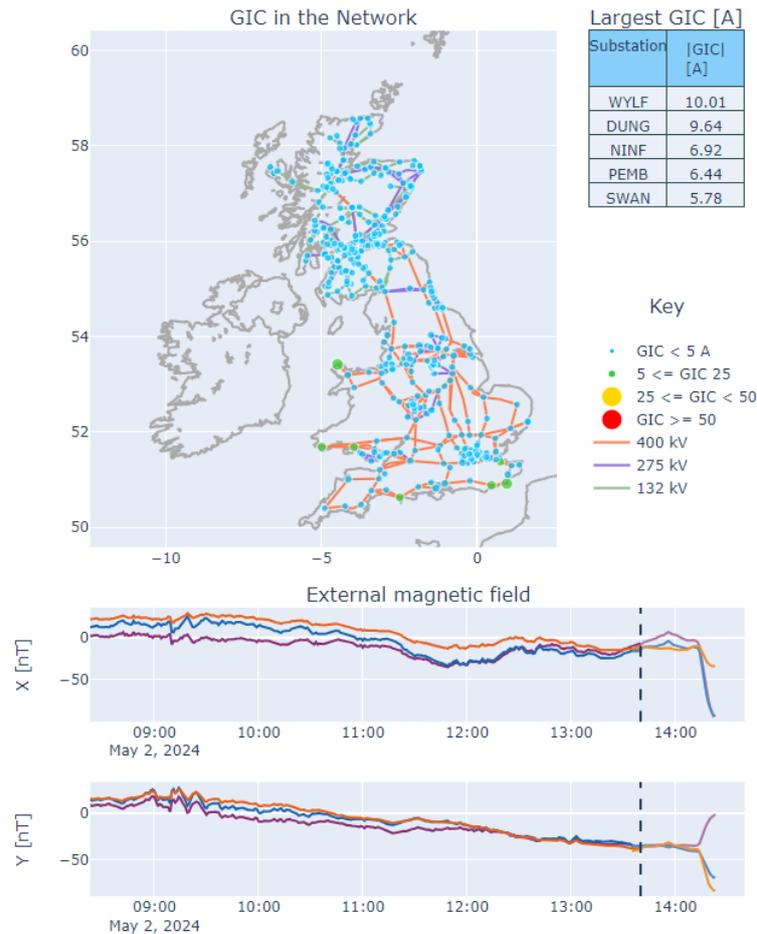


Figure 13: Similar to Figure 11, but for a GorgonOps GIC forecast. (upper panel) Forecast of GIC in the UK HV network using the magnetic field variation from GorgonOps. (lower panels) Magnetic field measurements (darker lines) and forecast from GorgonOps (coloured lines). Dashed vertical line indicates current time (note the addition of the previous 5 minutes of forecast data and the baseline normalization). Rail and PSP visualizations also mimic those produced by SPIDER. BGS © UKRI 2024

Substorm Forecast/M1: Substorm occurrence likelihood

The Substorm Forecast/M1 code forecasts the occurrence likelihood of a substorm in the next 1 hr period (depending on solar wind velocity). It is based on a Convolution Neural Network applied to L1 data that predicts substorm onset likelihood within next hour.

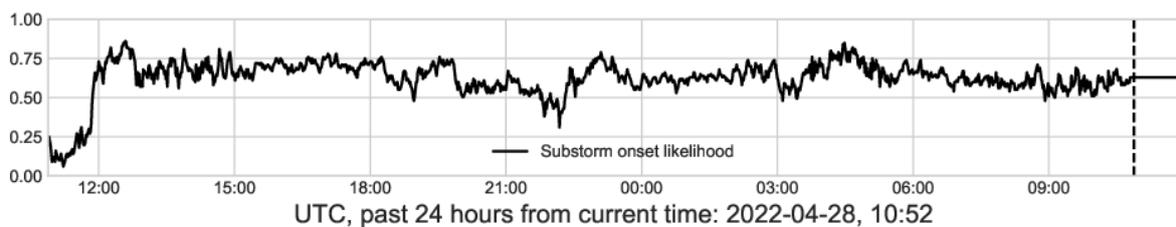


Figure 14: Example of the forecast occurrence of a substorm based on solar wind parameters. Dotted line indicates the forecast period. BAS © UKRI 2024

Figure 14 shows a time-series of likelihood of a substorm occurring during an active geomagnetic period in April 2022. The forecast period is indicated to the right of the dotted line with a suggested value of ~ 0.66 .

Shock Impact Assessment/M2 forecast: storm sudden commencement

The Shock Impact Assessment/M2 code is an optimised machine learning model applied to L1 solar wind data predicts storm sudden commencement (SSC) within next 4 hours. It predicts the largest dB/dt expected to occur in the UK within three days based on the magnitude of solar wind pulses that leads to geomagnetic storms.

The model code ingests L1 solar wind data and predicts the probability a sudden rapid change of the magnetic field based on analysis of magnetic field variation from the three UK observatories from 1997 to 2020.

An example of the probability of a shock in the solar wind and then a consequential sudden storm commencement (SSC) is shown in Figure 15. The two upper panels provide the L1 solar wind parameters over 27 hours, with the bottom panel showing the SymH index for comparison. The Shock Impact Assessment model probabilities are shown in the third and fourth panel showing the estimated probabilities of an SSC occurring (after the solar wind shock seen at L1). If the probability output by the shock identification model is greater than 0.7 then it is 'accepted' that there is a shock and the point is highlighted in orange in the SSC model (fourth panel).

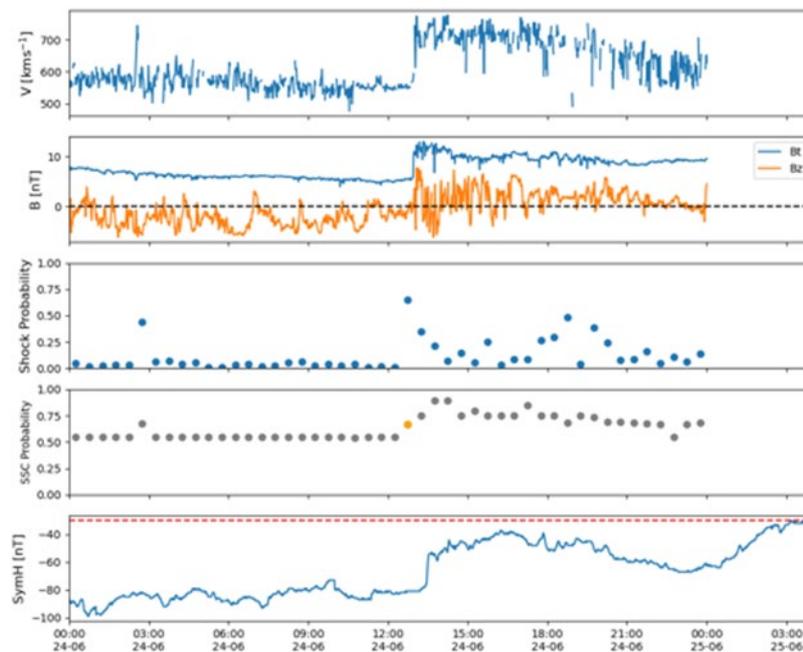


Figure 15: Example of the sudden storm occurrence forecast (fourth panel) along with L1 solar wind velocity (upper panel), IMF total field (t) and vertical (z) strength (second panel) and the probability of a shock in the solar wind (third panel). The lower panel shows the Sym H index. © MSSL 2022

Extreme Threshold Exceedance Forecast/M3 forecast: Exceedance of threshold rate of change at UK observatories

The Extreme Threshold Exceedance Forecast/M3 model is Recurrent Neural Network coupled to L1 (and/or MHD model) outputs and predicts the likelihood of rate of change of the magnetic field (dB/dt) being above a set of thresholds at each UK observatory. Again, this is based on training of the neural network against the rate of change of the magnetic field recorded at each

of the UK observatories. The thresholds are set based on the 90/95/99% occurrence in the observatory magnetic dataset. The probability of exceedance is then calculated from the relationship to the solar wind parameters.

There are several ways to visualise the data. Figure 16 illustrates two methods. The left column shows time series of the solar wind parameters (upper two panels) with the R parameter (rate of change of the field in nT/min) and the prediction from three different versions of the neural networks. The coloured lines are the threshold to be exceeded at a particular observatory during a storm on the 21st June 2015. The green line shows the probability of exceeding 18 nT/min.

The right column shows a snapshot of the probability of the exceedance of each threshold value at each observatory (LER, ESK and HAD). For a particular model run, the probability of exceedance of the magnetic field rate of change in the next 1 hour is shown as concentric circles.

At Lerwick (top), the chance of exceeding 18 nT/min (green circle) is 99% but exceeding 90 nT/min is 49%. For Hartland (bottom) the probability of exceeding 18 nT/min is 89% and 42 nT/min is 69% and so on. Note the correspondence between the thresholds is not perfect so inconsistencies can appear e.g. at LER, the red circle exceeds the orange one (when it should be lower).

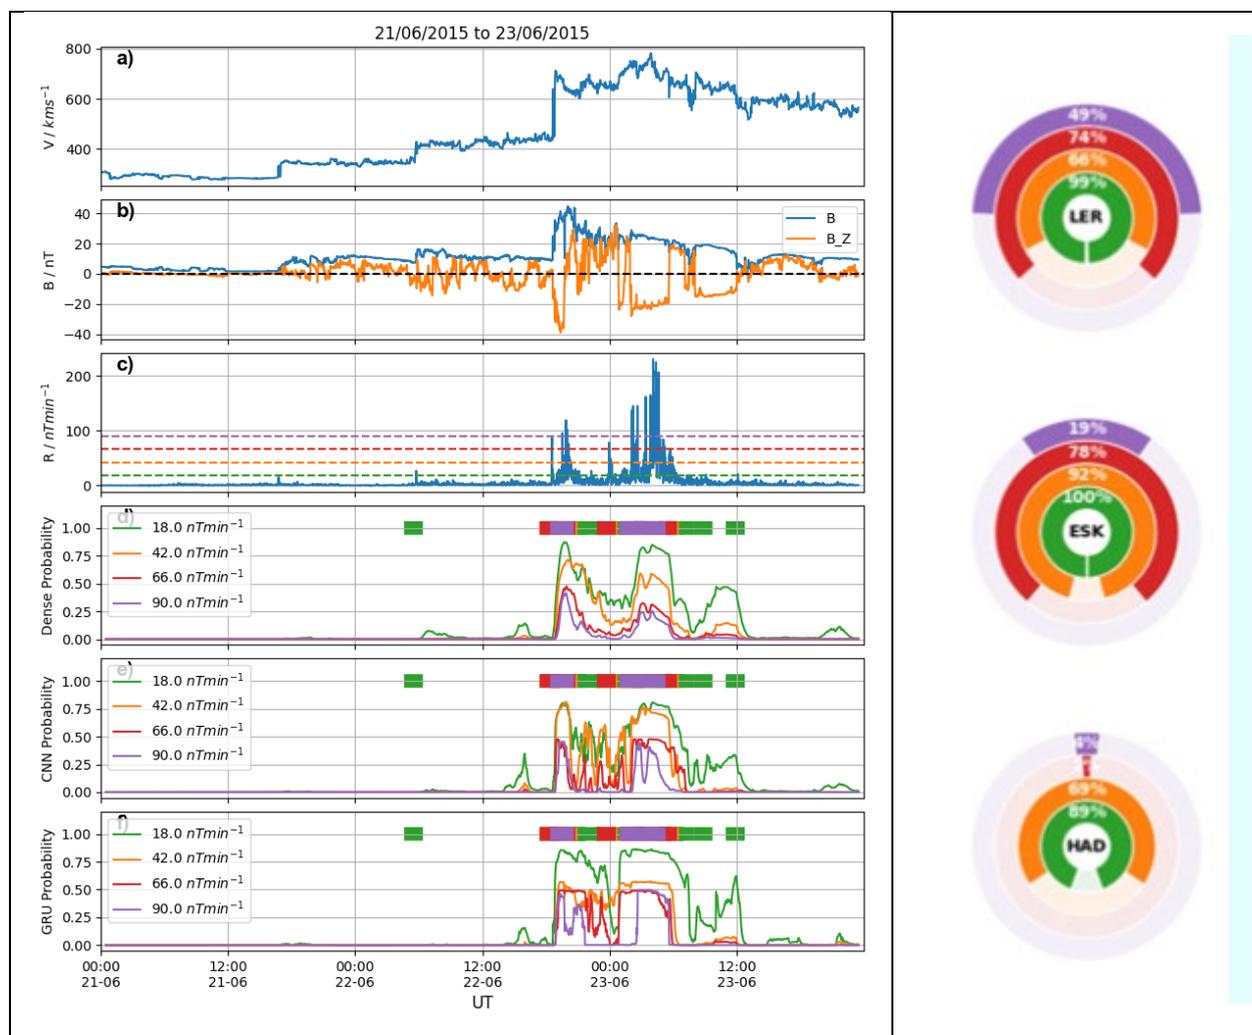


Figure 16: Example of the exceedance of the rate of change of the magnetic field. (left) As a time series showing L1 solar, predicted rate of change of the magnetic (R) and three separate outputs from various NN models. (right) Visualisation of a snapshot showing concentric circles of probability exceedance. © MSSSL 2022

3.1.4.4 GORGONOPS VISUAL DATA

The GorgonOps model has additional capability to [provide visualisation](#) 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 (Figure 17) or the location of the auroral oval over the UK (Figure 18). These visualizations are produced every 5 minutes and all relevant time instances are included for given forecast timestep, i.e., the full forecast horizon. Note that the timestamp in the images is that of the simulation time, whereas the file names indicate the actual forecast time.

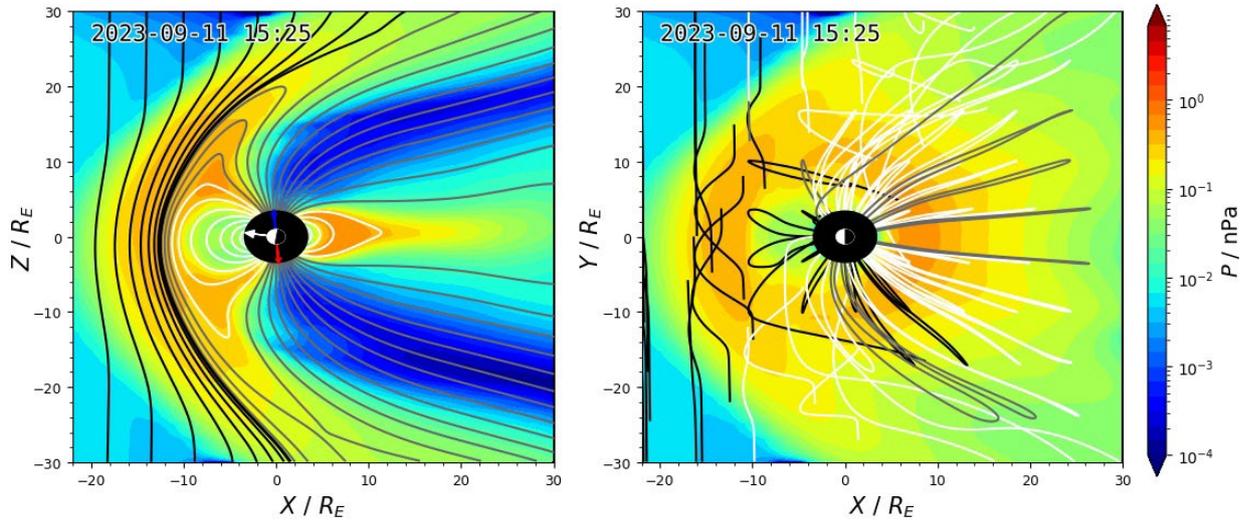


Figure 17: GorgonOps estimate of dynamic pressure (in nPascal) within the magnetosphere. © Imperial College 2024

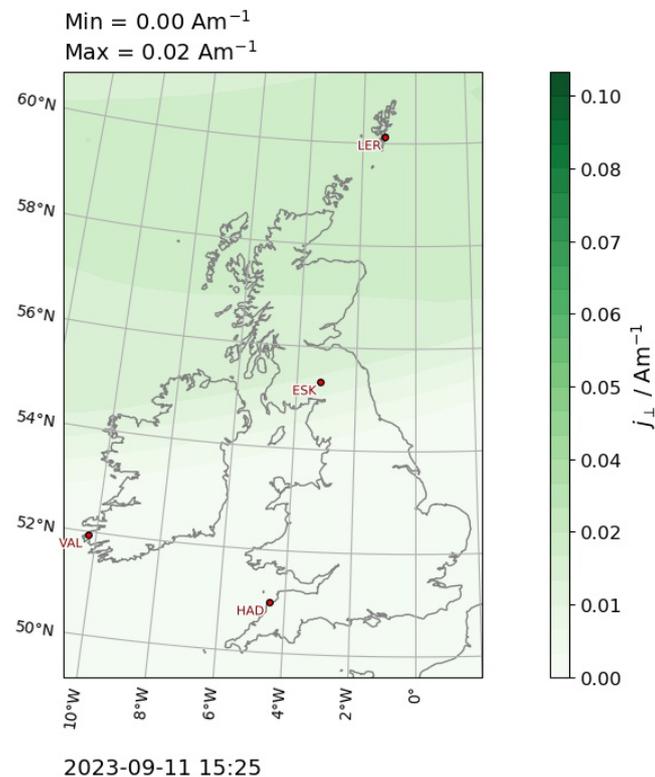


Figure 18: GorgonOps estimate of auroral current density over the UK. © Imperial College 2024

3.2 CONTAINERS

The following section describes the inputs, behaviour and outputs from the point of view of each Docker container in the <https://github.com/MetOffice/swx-swimmr-n4> repository. Figure 19 shows the flow of data through the system and outputs produced.

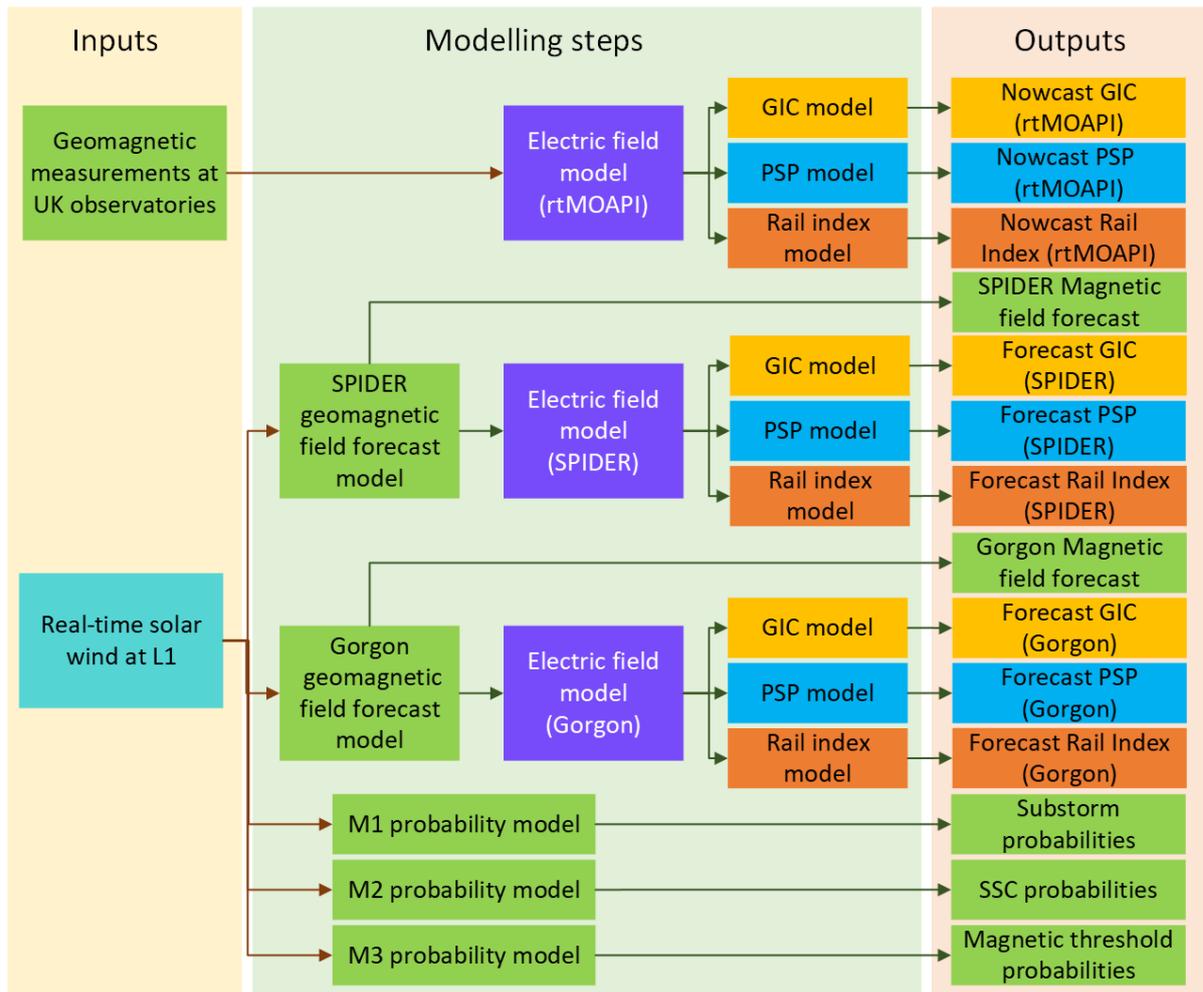


Figure 19 Flow diagram showing all the different inputs and outputs in the SAGE project. BGS © UKRI 2024

3.2.1 British Antarctic Survey Containers

3.2.1.1 SPIDER/SUBSTORM FORECAST (N4-MODEL-BAS-SPIDER)

There is only one container for the SPIDER/M1 docker (*n4-model-bas-spider*). It contains the machine learning training and forecast code for real-time predictions of the UK ground geomagnetic field, and substorm onset likelihood. It has four primary programs which underpin the SWIMMR-SAGE modelling.

The training programs are:

- *GGF_Training_Model.py*: This program trains the ground geomagnetic field (GGF) forecast models.
- *Substorm_Training_Model.py*: This program trains the substorm onset likelihood forecast model.

However, the models are already trained so this code is not used (but is available for updates later).

The forecasting programs are:

- *GGF_RTF.py*: This program forecasts ground geomagnetic field in real time, for about an hour into the future

- *GGF_RTFH.py*: This program forecasts ground geomagnetic field in real time, for about an hour into the future, and provides a hindcast for 24 hours into the past.

More detailed descriptions of each of these programs is found within the documents in the docker (e.g. header comments) of the code files themselves. These metadata descriptions also describe how to run the programs.

The forecasting code outputs are placed into an ASCII file called:

real_time_magnetic_field_forecast_from_program_GGF_RTFH_version_BRTFv2p2.dat

which is an ascii-formatted text file of the forecast values from the regression model ensemble (for the ground geomagnetic perturbation) and the convolutional neural network model (for the substorm onset likelihood). The forecast values extend about one hour into the future; the exact forecast span is dynamic and depends on the solar wind speed.

The Substorm Forecast model is the Substorm onset probability value which is between 0 (no likelihood of a substorm onset) to 1 which indicates a substorm is imminent in the next 1 hour period.

The SPIDER model is currently run every 5 minutes with the magnetic field outputs (the mean value) used to generate the geoelectric field.

One issue to note is that the SPIDER model requires two hours of *uninterrupted* solar wind values so often cannot run if there are missing values within the API due to data gaps in the L1 satellite download. It is also the case that there are long missing periods of data from the solar wind missions in the operational data (See Section 5.4).

3.2.2 British Geological Survey Containers

3.2.2.1 E-FIELD (N4-MODEL-BGS-EFIELD)

The [Efield outputs](#) are generated by the *calculate_electric_field.py* code which uses BGS *bgs_op_giclib* libraries to read magnetic field values from the Met Office API or the SPIDER or GorgonOps outputs. The filepaths are set in *file_paths_settings_efield.py* code.

The thin-sheet model FORTRAN code is stored in *progs/* folder and compiled during docker build. The conductance model (*conductance300902model_3.xyz*) and 1D resistance model (*resistmodels.txt*) are available at the top level.

The *efield_cron* initiates the run of the model every 5 minutes.

There are some simple unit tests in the *tests* directory.

The outputs are placed into the */root/efield/E_computed_{MODE}.dat* where mode is nowcast or a forecast from SPIDER or GorgonOps.

3.2.2.2 PIPELINE (N4-MODEL-BGS-PIPELINE)

The *psp_processing.py* reads in the output map generated by the *calculate_electric_field.py* code and also uses BGS *gicpylib* libraries to read in */root/efield/E_computed_{MODE}.dat*.

The [docker](#) contains the data files (csv and npz) for computing pipe to soil potential in the high pressure gas network. It also requires *pipelinelib* (a BGS library) and *bgs_op_giclib*.

The code runs once the *run_psp_code_watchdog.py* detects a change in the *success_{MODE}.txt* file written out by the Efield docker.

The outputs are placed into

/root/efield/psp_results/n4_bgs_psp_{MODE}/year/month/day/hour/minute/pipeline_psp_{MODE}.csv

A small python file [plot_pipeline_results.py](#) can output visualisation of the results (e.g. Section 0).

3.2.2.3 RAIL (N4-MODEL-BGS-RAIL)

The *rail_processing.py* reads in the output map generated by the *calculate_electric_field.py* code and also uses BGS *gicpylib* libraries to read in */root/efield/E_computed_{MODE}.dat*.

The *docker* contains a single data file (*UKrailnetwork.json*) for computing the rail index in various sections of the network. It also requires *railindexlib* and *pipelinelib* (BGS libraries), as well as *bgs_op_giclib*.

The code runs once the *run_rail_code_watchdog.py* detects a change in the *SUCCESS_success_{MODE}.txt* file written out by the Efield docker. The outputs are placed into */root/efield/rail_results/n4_bgs_rails_{MODE}/year/month/day/hour/minute/rail_indicator_{MODE}.csv*

There is a function in the python file *plot_rail_results* to allow visualisation of the results (e.g. Section 0).

3.2.2.4 POWER GRID (N4-MODEL-BGS-GIC)

The *gic_processing.py* reads in the output map generated by the *calculate_electric_field.py* code and also uses BGS *gicpylib* libraries to read in */root/efield/E_computed_{MODE}.dat*.

The *docker* contains a single data file (*UK_2022_earth0.5.pkl*) holding the representation of the UK HV power network in *pickle* format.

The code runs once the *run_gic_code_watchdog.py* detects a change in the *success_{MODE}.txt* file written out by the Efield docker.

The outputs are placed into */root/efield/gic_results/n4_bgs_gic_{MODE}/year/month/day/hour/minute/gic_{MODE}_{date}.csv*

There are three outputs – substation GIC, line currents and a set of plots

There is a function in the python file *plot_gic_results* to allow visualisation of the results.

3.2.2.5 GENERAL (N4-UTILITY-BGS-GENERAL)

The BGS *general container* is the only “utility” container in the model. It copies and deletes information from the outputs of the other BGS containers. It contains two shell scripts *copy_script.sh* and *tidy_script.sh*.

The copy script runs every five minutes and copies the outputs from the dockers into the */output/* directory for Met Office API to collect.

The tidy script runs hourly and deletes data files older than six hours from the volumes to control the amount of memory being used.

The reason for the general container is because the contents of the */output* folder are cleared every 5 minutes. The utility container creates a one-way copy which prevents the model outputs from being cleared from the containers they belong to.

We note another solution would be to use a base image for all the model images that copies output from an environment variable to */output*. That environment variable could be set to the output directory of each of the constituent models. This has the upside that it reduces a single point of failure, but it is also more complex and less flexible.

In future, other utility containers could be created from the utility image to perform tasks such as monitoring and validation.

3.2.3 Imperial College London Containers

3.2.3.1 MONSOON UNPACK

This [container](#) unzips the GorgonOps model output from its binary format and unpacks the data into a readable structure, accessible to both other containers and to the Met Office S3 bucket. The *monsoon_unpack.py* script additionally ensures that the unpacked data is shifted to its appropriate forecast horizon and includes the baseline normalization needed for ground geomagnetic field estimates at the UK observatories. On AWS, the *n4-model-icl-monsoon-unpack* container runs the python script every 5 minutes using a cron job, unpacking the latest forecast ground geomagnetic field and making the data available to both downstream containers in the appropriate output file structure (root being /output/gorgon). Logging is included both the AWS and Monsoon environments, with the latter packaged along with the output data. On the Monsoon2 HPC environment outputs are packaged, zipped and transferred to AWS, where there are unpacked to the *monsoon_unpack_out* volume, which is accessible to the downstream containers. The unpacked data is similarly to other models copied to the *copy_out* volume using a copy script, which is mapped to the SWIMMR output/ file structure and is available via the S3 bucket. The residual data is cleaned out every 12 hours, whereas the persistent data is cleaned every 2 hours (should already be backed up to the S3 bucket by this point).

3.2.3.2 GORGON-TOOLS

This [docker](#) contains tools for extracting the relevant data from within the GorgonOps model output, with additional tools to compute derived parameters and produce bespoke visualisations (currently all done on the Monsoon HPC environment). Some functionality that may be incorporated in the future is the use of the Complex-Image Method (CIM) to compute the ground geomagnetic field with both internal and external contributions. This method was developed as part of the SAGE project, but validation efforts showed that the line of sight Biot-Savart integration done inline is more accurate and more efficient.

3.2.4 University College London Containers

3.2.4.1 SHOCK IMPACT ASSESSMENT (N4-MODEL-UCL-M2)

The [sudden storm commencement](#) shock probability is computed using the *M2_ModelRun.py* code. The code accesses L1 solar wind data using the API, interpolates to fill in any missing values up to 15 minutes long. It then loads in a series of machine learning models (*Models/Shock_ID_*) and features (*Models/scaler_3Param.p*) to compute the shock probability.

The output is written to *Outputs/M2_{date}.csv*

3.2.4.2 EXTREME THRESHOLD EXCEEDANCE FORECAST (N4-MODEL-UCL-M3)

The Extreme Threshold Exceedance Forecast/[M3 code](#) runs in a similar manner to Shock Impact Assessment/M2 docker. The *M3_ModelRun.py* code collects L1 data from the API and process it using a convolutional neural network (CNN) to predict the probability of exceeding a set of dB/dt thresholds at the UK magnetometer sites, based on a 12 previously trained CNN models in the *Models* directory within the docker.

There is a series of visualisation functions to plot out the examples in Figure 16.

The output is written to *Outputs/M3_{date}.csv*

3.3 OPERATIONAL SEQUENCING

The SWIMMR API updates and exposes magnetic and solar wind data every minute. Most of the SAGE models require a few minutes to a few hours of data in order to produce an output. There is a trade-off with processing time and update rate of the model. Processing time is on the order of one to two minutes depending on the model, so that is the update limit. Updates of 30 minutes would be too long and 2 minutes is too fast. The system is therefore currently set to update every five minutes.

The code processing is controlled by *crontab* in each of the dockers and is set to run the main processing code every five minutes, assuming certain conditions are met.

3.3.1 Nowcast

The E-field code in nowcast mode triggers every 5 minutes and reads the BGS data from the API. It computes the geoelectric field map 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 1 minute to ascertain whether the *success* file's timestamp has changed. When it does, this triggers them to compute their respective outputs. Figure 20 illustrates the sequence.

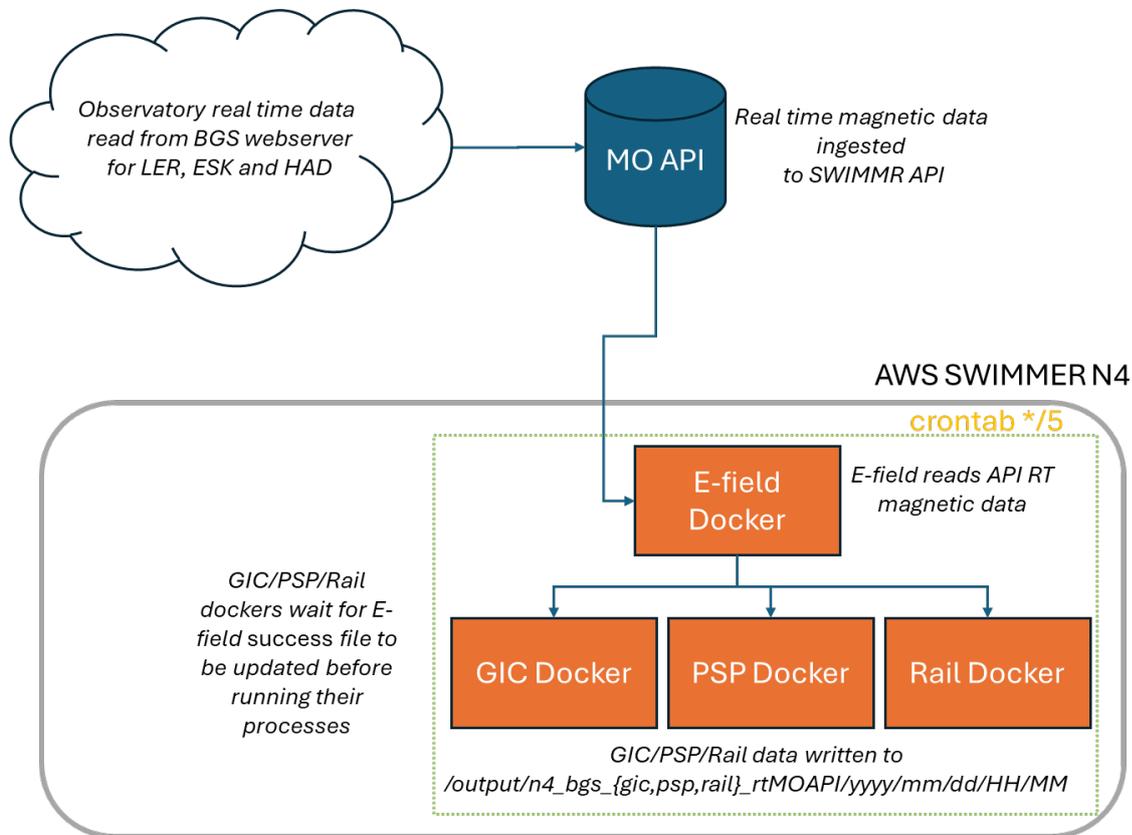


Figure 20: Sequence of processing for the nowcast models. BGS © UKRI 2024

3.3.2 Machine Learning Forecast

For forecast mode, the SPIDER model is run every 5 minutes, again controlled by the *crontab* 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 code in forecast mode is triggered every 5 minutes and reads the SPIDER data file from the docker-compose environment if a new file is detected.

It computes the geoelectric field map 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 1 minute to ascertain whether the *success* file's timestamp has changed. When it does, this triggers them to compute their respective outputs. Figure 21 shows the sequence.

Similarly, the Shock Impact Assessment and Extreme Threshold Exceedance Forecast models are triggered to run every 5 minutes using *crontab*.

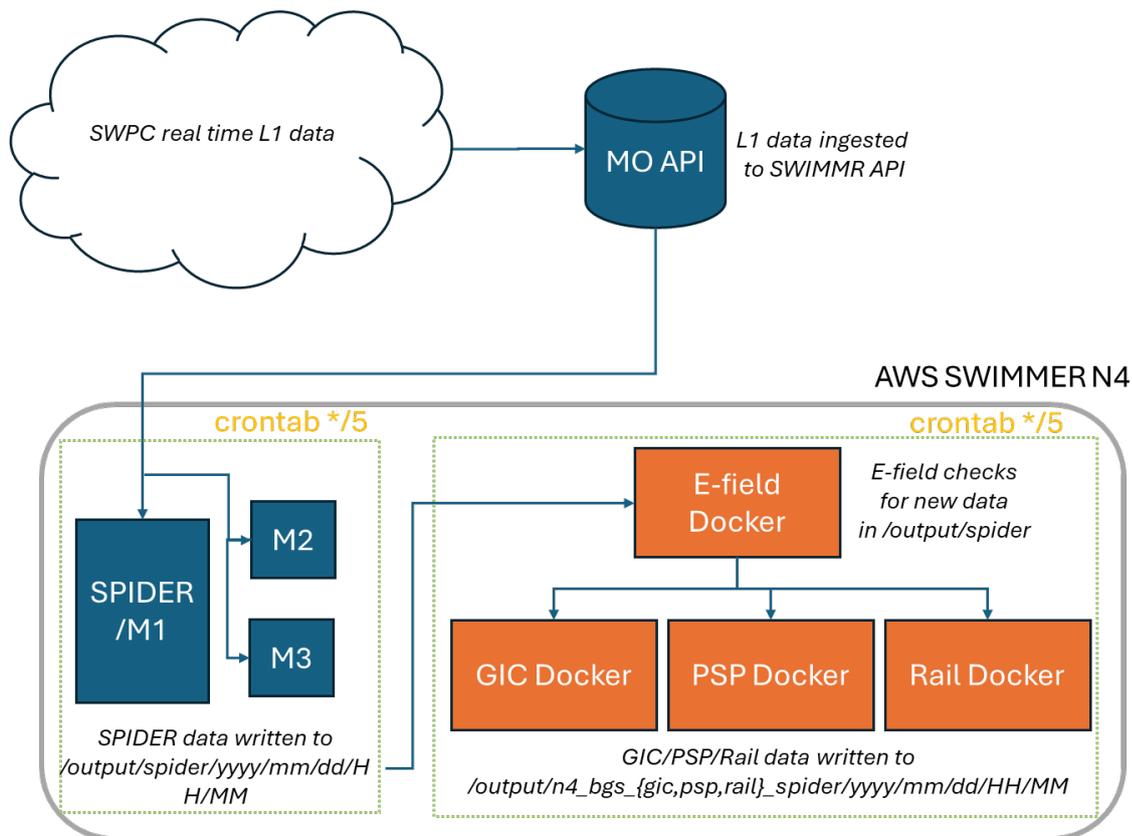


Figure 21: Sequence of processing for the machine learning forecast models. BGS © UKRI 2024

3.3.3 GorgonOps Forecast

For GorgonOps model, this runs on the Met Office HPC (presently Monsoon) due to the larger computational requirements. The model is run continuously, producing output at a minute cadence. The ingestion cycle for real-time solar wind input on Monsoon is every 5 minutes, which is linked to a matched transfer cadence of every 5 minutes when a processing script run in parallel uploads a compressed file (zipped) to the API. On AWS the `monsoon-unpack` docker unpacks and processes this output. The docker is controlled by a `crontab` file, which also runs on a matched 5-minute cadence. All processing in these cases is greedy, making use of all available data for processing at the point of being triggered. Each iteration of the cron job generates a new file with the magnetic field forecast for the following hour, stored in the appropriate directory structure. The E-field code in forecast mode is triggered every 5 minutes and reads the GorgonOps magnetic field forecast data file from the docker-compose environment.

It computes the geoelectric field map 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 1 minute to ascertain whether the `success` file's timestamp has changed. When it does, this triggers them to compute their respective outputs. Figure 22 shows the sequence of processing.

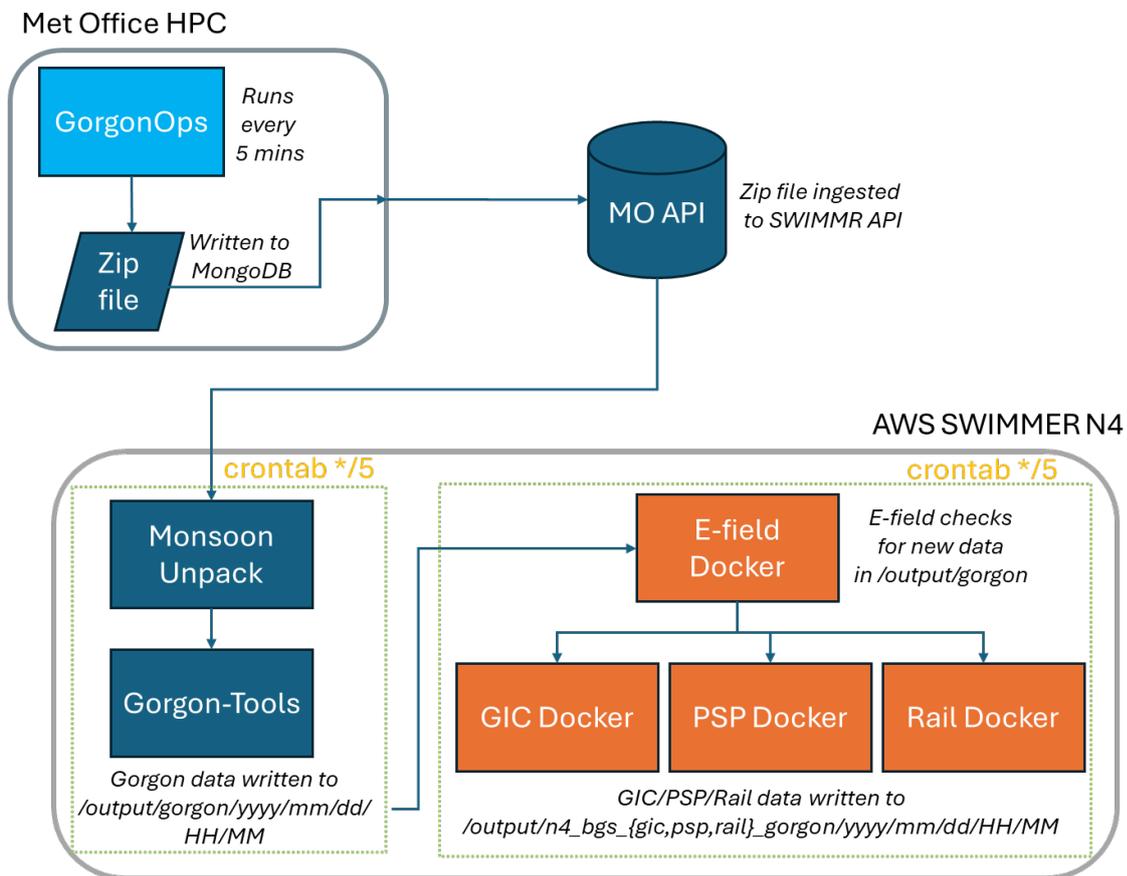


Figure 22: Sequence of processing for the GorgonOps forecast models. BGS © UKRI 2024

4 Guide for Developers

4.1 REPOSITORY

The model code is hosted on GitHub in the private project [swx-swimmr-n4](#).

(Within BGS, a shared e-mail exists so that BGS staff can access the repository using bgs.sage.spaceweather@gmail.com.)

Authenticating `git` actions that manipulate the GitHub project, such as pulling and pushing requires more than just the username and password. You must first set up a personal access token. You can do this by going to the [personal access tokens page](#) (If the link doesn't work, from any page on GitHub follow the path: sagebgs profile icon → settings → developer settings → personal access tokens). You then use the token when interacting with the repository instead of the account password.

It is good practice to have a unique branch for new features or bug-fixes, merging into the main branch (in this case the main branch is called 'develop') when the changes are agreed to be stable and beneficial. Currently, merging into develop and then pushing develop to origin (the GitHub repository) triggers a build within the Met Office AWS ecosystem of all the images defined in the Docker compose file. More documentation about the build environment can be found in the [Met Office Building a model documentation](#).

The Docker image names are defined as environment variables in the `.env` file within the repository. Note the very same environment variable names are also made available in the build environment within the AWS ecosystem (if environment variables are defined in the `.env` file, the values assigned in the file take priority). If new images are required in the build step (for example, if another model component is added to SWIMMR N4) they will not work in the AWS ecosystem until a space has been made available in the container registry for them. New spaces can be added by opening a dialogue with the Met Office at SWIMMR@metoffice.gov.uk.

4.2 DEPLOYMENT

A successful (or unsuccessful) build after a merge into origin/develop will send out e-mail alerts to the e-mail addresses defined in `pipeline-contacts.json`. The build does not automatically trigger creation of the model containers, but it does cancel the existing model run. To restart SWIMMR N4, follow these steps from a bash console:

1. Initialise an SSH agent with: `eval `ssh-agent -s``
2. Add a key³ to the agent with: `ssh-add sage.pem`
3. Login to a central 'bastion' machine with:
`ssh -i sage.pem -A sage-user@bastion.swimmr.spaceweather.metoffice.cloud`
4. Access the execution environment with:
`ssh ec2-user@10.0.148.154`
5. Trigger the model with `sudo trigger_model`

The `trigger_model` command is the only command that `ec2-user` is allowed to run with root privileges, no password is required. If the model is running successfully, outputs and logs will begin to appear in the [portal](#) and will be accessible from the [Met Office API](#). More information on deployment and the execution environment (such as how to configure the `execution-config.json` file) can be found in the [Met Office documentation for Deploying a Model](#) and the [documentation for the Execution Environment](#). Certain Docker objects, such as volumes can persist across model versions and some changes to the model might require manual intervention.

³ This requires finding the relevant key/asking the Met Office or a colleague for it.

4.3 LOCAL DEVELOPMENT

4.3.1 Overview

It is a very slow feedback process to add features to one of the constituent parts of the model, get the changes approved so they can be merged into the develop branch, and only then see if the resultant model runs in the AWS ecosystem. To that end, this section documents how to run SWIMMR N4 on local infrastructure for debugging purposes. The idea is that if the user has access to Docker, they can emulate the AWS environment and perform build and run steps in a nearly identical virtual environment but with much more access and control over the model.

4.3.2 Replicating the AWS Environment

Included in the repository is a Dockerfile simply called 'Dockerfile'. This is the Dockerfile that ultimately will create an Amazon Linux 2 container with the required software mentioned in Software and Infrastructure – Docker and Docker Compose. This leads to a [Docker in Docker](#) solution. [Docker in Docker is generally not recommended](#) but this setup is one of the most accurate ways to reflect the AWS environment the model would run in, so its use is somewhat justified. Improvements that could be made to improve the security of SWIMMR N4 are detailed in Security.

The first step to replicating the AWS environment is to clone [swx-swimmr-n4](#):

```
git clone https://github.com/MetOffice/swx-swimmr-n4.git
```

```
cd swx-swimmr-n4
```

There are some environment variables set in the AWS (probably using the secret 'environment-secrets') that need to be replicated. In the root of swx-swimmr-n4, create a file 'ENV_VAR.env' that contains the following environment variables:

```
GICUSER=<put-gic-library-user-here>
GICPW=<put-gic-library-password-here>
API_KEY=<put-met-office-private-api-key-here>
```

and create another file called 'secrets.env'⁴ which sets the following environment variables:

```
BGS_OP_GICLIB_USER=<put-giclib-library-username-here>
BGS_OP_GICLIB_PASSWORD=<put-giclib-library-password-here>
PIPE_USER=<put-pipeline-library-username-here>
PIPE_PASS=<put-pipeline-library-password-here>
RAIL_USER==<put-rail-library-username-here>
RAIL_PASS==<put-rail-library-password-here>
```

Now you can build an image which replicates the AWS environment. Build the image with the command:

```
docker build -t <name-your-image> .
```

There are several different types of containers you might want to run from the image. The most useful for development purposes is to mount your swx-simmr-n4 folder inside the container. This makes it so any changes you make locally are reflected in the new environment and is controlled by the -v flag. It is recommended to run the container in the background (-d) because to begin with the container will just log the output of the Docker server it is set to run on startup. Finally, the --privileged flag enables a container to run Docker in Docker, so that the final command is:

```
docker run -d --privileged -v <path-to-swx-swimmr-n4>:/home/ --name
<name_your_container> <name-of-image>
```

To access the newly created container, use [Docker exec](#) with the interactive flag (-it):

⁴ You can name this file anything you like so long as you are consistent throughout development.

```
docker exec -it <name-of-container> bash
```

```
cd /home
```

The very first time the bash command is run inside your new container, set up a secrets file. To do that, run these two commands:

```
docker swarm init
```

```
docker secret create environment-secrets secrets.env
```

This is now a functional development environment. The model should now build and run with the command:

```
docker compose up
```

or

```
docker compose up <component-of-SWIMMR-N4>
```

to run only a specific part of the model defined in the docker-compose.yml file, for example 'spider'. The outputs from the model will be available in /output.

4.3.3 Maintaining the Development Environment

Unless further `-v` mounts are made from the development environment to the model component containers (which can be a good idea), every time a change is made to a model component, run `docker compose up` again. However, Docker compose does not rebuild images by default, so it may be desirable to rebuild the required images with the `docker compose build` command or remove the existing image with `docker rmi <image-name>`. Other items manually such as containers and Docker volumes can be managed too.

All the Docker containers take up disk space, so there are a variety of Docker commands to reclaim space, a good starting place is `docker system prune` which can be run from the development environment/container or the host docker machine.

4.4 FUTURE IMPROVEMENTS

4.4.1 Security

There are a few security improvements for `swx-swimmr-n4` that should be considered. Of most concern is vulnerabilities that could occur on the Linux systems as they host a huge range of valuable resources.

The biggest concern is during development, when the `--privileged` flag is used to run the Docker in Docker environment (detailed in [Replicating the AWS Environment](#).) That flag effectively allows the container root access to the host system. This is compounded by the fact that none of the images in the model set the container users to be anything other than root, which has been exploited in the past even without the `--privileged` flag to gain access to the host system as root.

The latter problem can be solved quite easily using the `USER` attribute inside all of the model Dockerfiles, guidance for which can be found in the [best practice for Docker guide](#). The former problem is more complex to solve. The most straightforward and general solution is to further separate the development and operations environments so that if something does go wrong during development the compromised environment cannot access BGS or personal data. There are many ways to go about this although they probably all involved a sandboxed virtual machine of some kind. Another improvement would be to start using [rootless Docker in Docker](#).

4.4.2 Resilience

There are some improvements that could maybe be made to increase the resilience of the model. In this context, resilience means how secure the model is against the failure. There are multiple modes of failure, for example the whole environment crashing all the way down to single containers.

Some primary modes of failure could be disk space or memory failures, where more and more files are written to the disk or variables get bigger and bigger upon each model iteration, so it is worth checking/testing for this. Then there are failures that could occur in operation because of strange/out of bounds inputs (0s, NaNs, bad API calls etc.). Testing and time will surface many of these bugs, and there is also the option to include [automatic container restarts in the Docker compose file](#).

The General container is particularly problematic, because if it fails then none of the BGS data will make it to the Met Office API output. It would be worth running tests to check its disk space usage is not increasing over time. Fortunately, it is also very simple (and should be kept that way) and therefore it should be unlikely to fail and easy to restart.

4.4.3 Geoelectric field model improvements

The next iteration of the geoelectric field computational code will move away from the thin-sheet modelling code and use the magnetotelluric-based space weather hazard model. This will require around 3 hours of magnetic field values prior to the current instant to produce a geoelectric field map. Development of this code is on-going at present.

5 Validation

Validation of models is important for monitoring behaviour and to establish a baseline for improvements. The SAGE models are often not easy to ground truth as the outputs are either not physical or cannot be readily compared against measurements (as they do not exist).

The products fall into five loose categories:

- 'feasible': can be verified directly, continuously, and automatically in AWS (SPIDER; GorgonOps; SWMF)
- 'possible offline': cannot be verified within the AWS as it needs expert judgement or data from an external third-party provider (substorm lists; GIC measured in the GB HV grid)
- 'not readily measured': geoelectric field modelled outside of the UK observatories
- 'not measured': GIC in pipelines
- 'not physical': the rail index

Table 1 summarises the products that can be verified and the 'ground truth' comparison data.

	Products that can be verified directly on AWS	Products that could be verified later (not on AWS)	Products that cannot be verified
Product	<p>SPIDER forecast of ground magnetic field</p> <p>GorgonOps forecast of ground magnetic field</p> <p>SWMF forecast of ground magnetic field</p>	<p>Substorm Forecast (M1): Forecast of substorm occurrence</p> <p>Shock Impact Assessment (M2)/Extreme Threshold Exceedance Forecast (M3): SSC & Forecast dB/dt</p> <p>Nowcast/forecast of GICs in power grid</p>	<p>Nowcast/Forecast of pipeline PSP</p> <p>Nowcast/Forecast of rail index</p>
'Ground truth'	<p>BGS real-time observatory data</p>	<p>Post-event list of manually identified substorms and dB/dt</p> <p>Post-event comparison to GIC measurements at 2 UK sites</p>	<p>Not presently measured;</p> <p>Not physically realised</p>

Table 1: Summary of products and their verification potential.

5.1 PRODUCTS THAT CANNOT BE VERIFIED

There are two products that cannot be verified as either there are no measurements made in the UK.

5.1.1 Nowcast/Forecast of pipeline Pipe-to-Soil Potential

There is ample evidence of GIC flow in high pressure pipelines from measurements in pipelines around the world in Alaska, Finland and New Zealand (e.g. Ingham and Rogers, 2018). The pipe to soil potential (PSP) can be modelled successfully using a similar method to modelling the flow of GIC in the high voltage power network with the admittance representation (e.g. Boteler, 2013).

However, measurements of the PSP in the UK pipeline network are not available, despite being monitored by pipeline operators. Unlike line GIC, no proxy measurements of PSP are available either. Hence, the values of PSP predicted are entirely theoretical and not validated. They are intended to be indicative of where there might be problems experienced by operators, though presently there is little engagement with the industry so far.

5.1.2 Nowcast/Forecast of rail index

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 voltages (not current) can be modelled using a similar method to modelling the flow of GIC in the high voltage power network with the admittance representation (e.g. Boteler, 2021).

However, like measurements of the PSP in the UK pipeline network, measurements of GIC in rails are not available, as this is a difficult measurement to make on live rail networks (primarily for safety reasons). Hence, the values of rail index predicted are entirely theoretical and not validated. Future work to improve the modelling of GIC directly may arise in future (e.g. Patterson et al. 2023).

5.2 PRODUCTS THAT ARE NOT CONTINUOUSLY VERIFIABLE

There are several products that have been verified in scientific studies but cannot be continuously verified within the AWS environment.

5.2.1 Geoelectric field modelling using thin-sheet

The geoelectric field based on the thin-sheet model produces $179 \times 131 = 23449$ (10 km)² cell blocks. Only two of these have continuous measurements within them – at the ESK and HAD observatories. Around 50 other sites have a magnetotelluric survey within their boundary. That implies that $(52/23449) \sim 0.2\%$ of the cells have been checked in the thin-sheet model.

Where there are data available, the model has been compared in scientific papers (see Beggan et al, 2021). They found the thin-sheet model has large uncertainties (of over 50%) and typically underestimate the geoelectric field at the observatories or at MT sites where measurements can be compared against the model. This is the primary reason for moving from the thin-sheet model to the magnetotelluric-based geoelectric field in order to allow continuous improvement of the conductivity model.

5.2.2 GIC modelling in the GB HV power grid

It is possible to compare the values of GIC in the HV grid with GIC measurements at two substation sites in Scotland (Torness and Strathaven). These are the only GIC measurements available in the GB network. BGS have occasional access (by email request) to the GIC measurements from Scottish Power but these tend to be collected after large geomagnetic storms. Scottish Power only store GIC data for around 4 weeks before deletion.

To validate the GIC in the HV network model elsewhere in the GB grid, a large-scale campaign was developed in the SWIGS project and run between 2018 and 2022 to make indirect measurements of the GIC flowing in power lines using two magnetometers. The results from 12 sites in Scotland, Wales and England show excellent correlation and match of the amplitude of the line GIC. Note the National Grid network changes very little between years so the HV model is essentially static.

Papers by Huebert et al (2020,2024) demonstrate that the BGS HV model along with MT geoelectric field modelling captures the measured line GIC values to within 10% of the true value (example shown in Figure 23 for November 2018 storm). Correlation and polynomial fit coefficients are shown in the bottom left corner of each panel.

Given the other uncertainties and missing parameters (substation grounding being the primary one) in the modelling, we consider the HV model to be validated for the purposes of the SWIMMR N4 nowcast and forecasting capability.

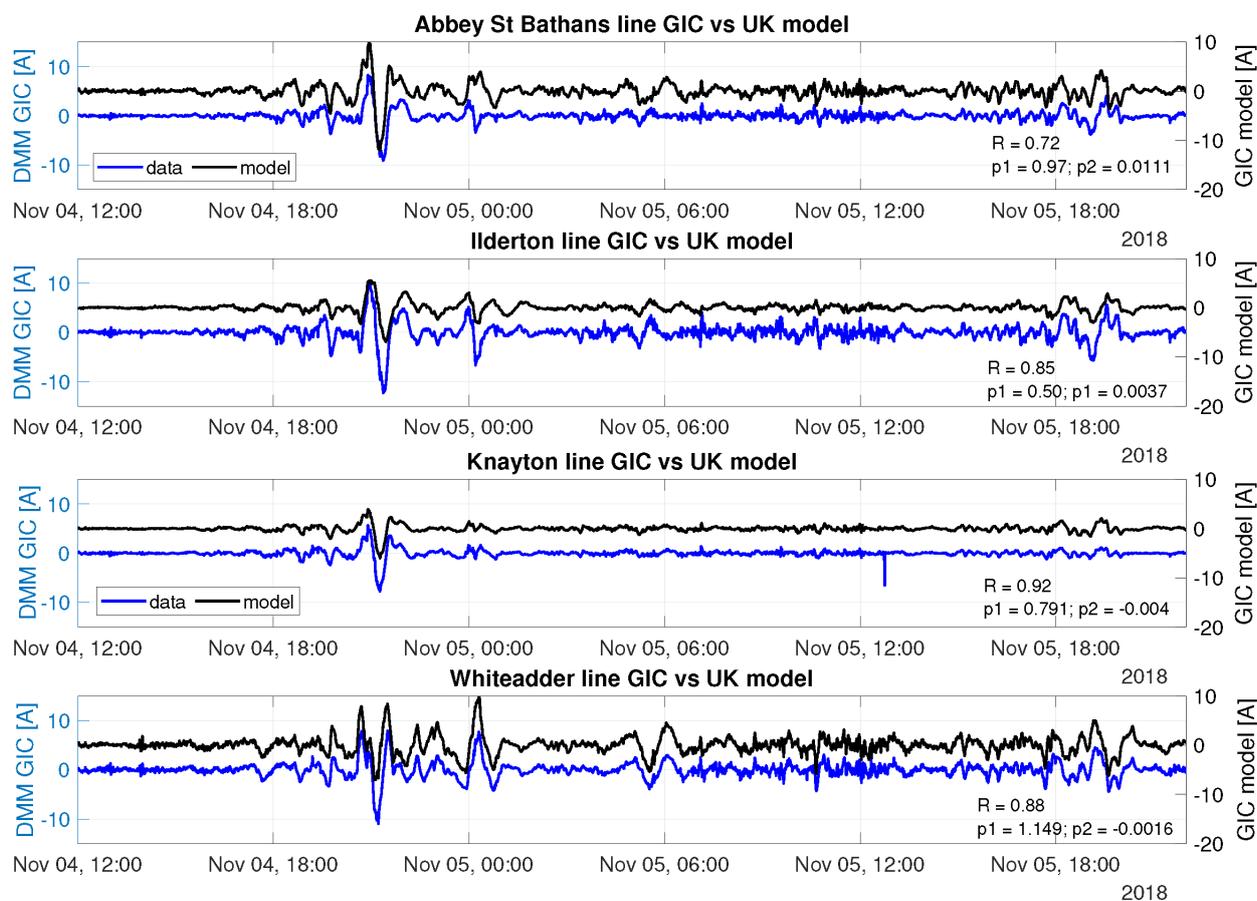


Figure 23: Measured (left axis, blue) and modelled (right axis, black) line GIC for sites Abbey St Bathans (ASB) [East Lothian], Ilderton (ILD) [Northumberland], Knayton (KNA) [North Yorkshire] and Whiteadder (WHI) [East Lothian] for the 4-6 Nov 2018. The zero level for data and model are offset for clarity. Correlation coefficient and polynomial fit coefficients are shown in the bottom left corner of each panel. BGS © UKRI 2021

5.2.3 Substorm probability modelling (Substorm Forecast)

The substorm machine learning model (Substorm Forecast) provides a regression relationship between the L1 solar wind parameters and the occurrence of substorms within a fixed lag time. The analysis is not, as yet, published so much of the detail resides in SWIMMR N4 partner or conference presentations. In essence, the modelling used a set of recognised substorms over a period of 18 years (1996-2014) to train the model to recreate the storms on the SOPHIE EPT90iso list (for events with waiting times between 1-6 hours).

It is difficult to determine how well the ML model actually predicts substorms because there is not a definitive list of substorm onsets (or even a good definition of what constitutes a substorm). The Substorm Forecast model provides a positive set of receiver-operator curves (ROC) when tested against substorms that it has not encountered before (post 2014). As a cross check, Figure 24 shows the ROC predictions against the lag time between the L1 parameters for different windows. The peak is for 0-2 hours which makes physical sense as this is the accepted response time of the magnetosphere and ionosphere to the forcing of the solar wind.

To repeat this evaluation for verification purposes would require a new list of substorms to be created in addition to the SOPHIE list (Forsyth et al, 2015) – a problem which requires expert manual judgement. The substorm list is not updated very often (e.g. every five years) and to verify the model requires a scientist-in-the-loop to run the model with the L1 data and compare it against the new substorm list.

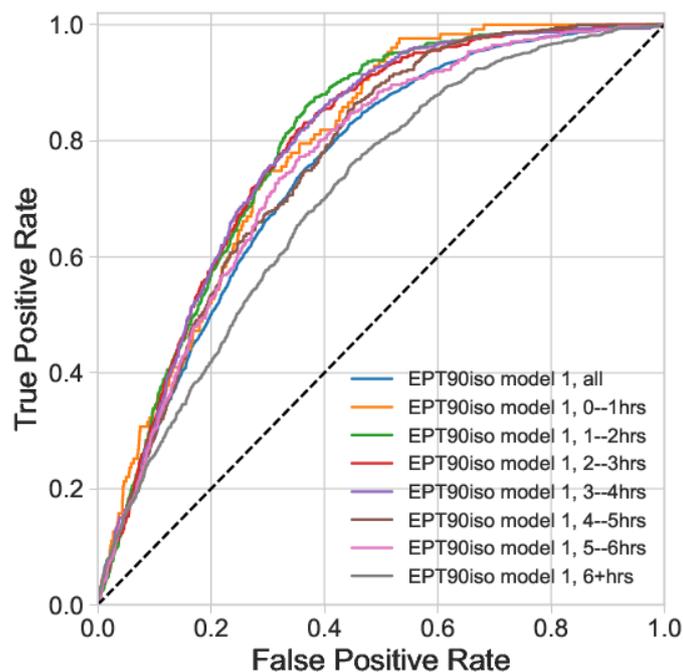


Figure 24: ROC curves from a model trained on the full EPT90 (substorm) list from 1996-2014, and curves computed from subsets of the EPT90 test set, for specific waiting times. BAS © UKRI 2022

5.3 PRODUCTS THAT CAN BE CONTINUOUSLY VERIFIED

The forecast magnetic field values from SPIDER and GorgonOps (and/or SWMF) can be compared directly to the magnetic field measured (after the forecast period) at the UK observatories.

5.3.1 SPIDER

The SPIDER code makes a *1 hour forecast* of the minute-mean magnetic field every 5 minutes for each observatory component at the three UK observatories (LER, ESK, HAD). This provides nine time-series to compare against. We evaluate the forecast magnetic field from SPIDER against BGS ground data after the hour has passed.

We can compute a variety of metrics including the root-mean-square (RMS) error in nT and the Pearson correlation (-1 to +1). If we set a threshold of ± 50 nT of the external field variation away from zero (which assumes a quiet day), we can compute whether the forecast has a hit (H) or miss (M) against the threshold value. From the number of hit and miss values we can compute the Heilke Skill Score (HSS), frequency bias, the Probability of Detection (POD) and Probability of False Detection (POFD). See Florczak et al (2023) and references therein for definitions.

The outputs of the forecast metrics are stored as an html visualisation and as csv files in relevant folder on AWS: `year/month/day/hour/min/`

As an example, Figure 25 shows the forecast and subsequent measured external field values for a G2 storm on the 18-Sep-2023 for 19:00-20:00 UT. The SPIDER model is the coloured line with the grey outline showing the 1σ uncertainty. The darker lines are the measured magnetic field in each component (x, y and z) at each observatory (esk, had and ler).

Table 2 shows the RMS and correlation metrics for the hour of the 18-Sep-2023 for 19:00-20:00 UT. If both the model and forecast continuously exceed ± 50 nT during the period, then the threshold metrics are computed. In this hour, only HAD Z meets this criterium. The Spider model performs relatively well in the additional metrics despite having low correlation.

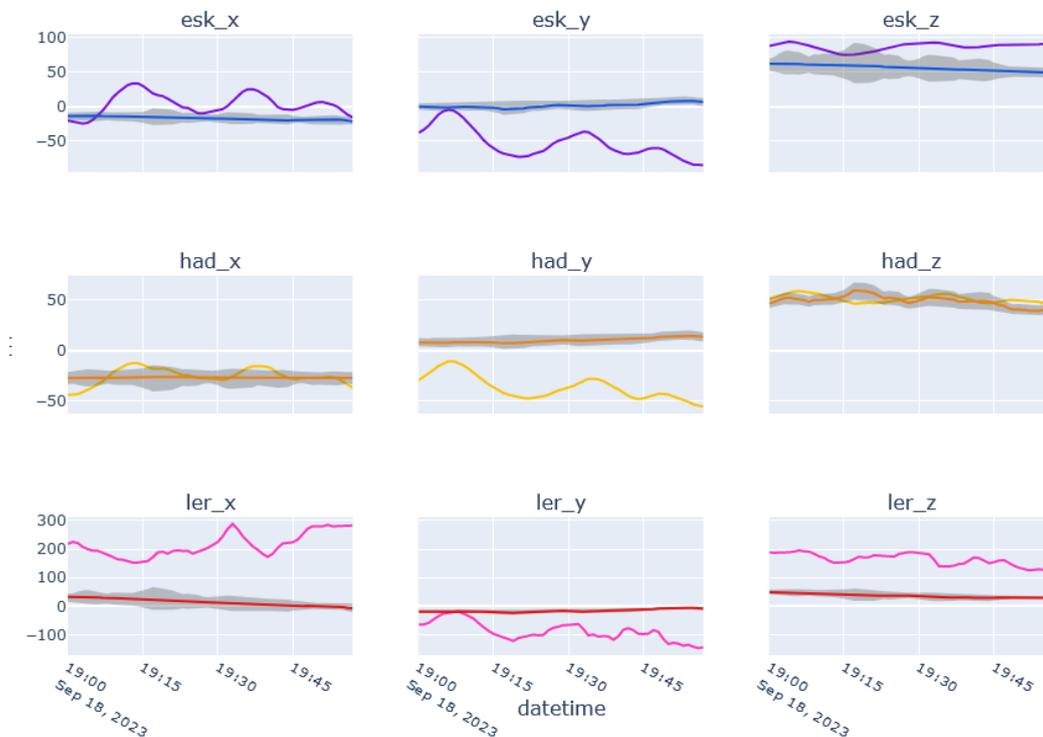


Figure 25: Comparison of the SPIDER model and uncertainty against measured data at the three UK observatories for the 18-Sep-2023 for 19:00-20:00 UT. BGS © UKRI 2024

OBS Component	RMSE (nT)	Correlation	HSS	FB	POD	POFD
esk_x	25.86	0.04	0.00	0.00	0.00	0.00
esk_y	59.17	-0.46	0.00	0.00	0.00	0.00
esk_z	31.84	-0.37	0.00	0.00	0.00	0.00
had_x	7.70	0.32	0.00	0.00	0.00	0.00
had_y	48.70	-0.60	0.00	0.00	0.00	0.00
had_z	6.24	0.10	0.19	1.03	0.63	0.44
ler_x	207.33	-0.70	0.00	0.00	0.00	0.00
ler_y	80.06	-0.53	0.00	0.00	0.00	0.00
ler_z	128.69	0.67	0.00	0.00	0.00	0.00

Table 2: Evaluation of metrics for 18-Sep-2023 for 19:00-20:00 UT.

In reality, determining what a good ‘fit’ is between the model and data is difficult. Florczak et al (2023) discuss this in detail noting that while the two time-series may look superficially similar (or ‘good enough’ for forecasting GIC), the metrics suggest otherwise.

The metrics for SPIDER are logged continuously when a model forecast is produced (see Section 5.4 for why this is not always the case).

5.3.2 GorgonOps

GorgonOps has been benchmarked according to scientific simulations following the procedure in Gordeev et al. The model has been optimised to reliably produce forecasts given various driving regimes and produce the required output (ground geomagnetic estimates at Lerwick, Eskdalemuir and Hartland) and visualizations (of the magnetosphere and auroral electrojet above the UK region) faster than real-time. (2015) (see Figure 26 and Figure 27). This benchmarking ensures that the operational configuration is able to reproduce trends seen in empirical models and measured data, as was done for full MHD models in the original paper. Further benchmarking relative to canned historic data has also been done, and is the focus of continued work.

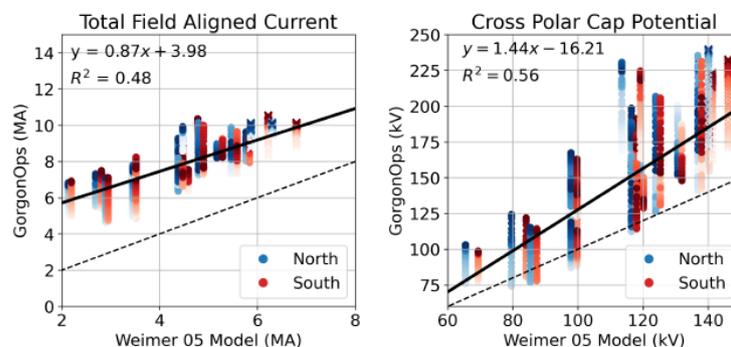


Figure 26: GorgonOps benchmarking relative to the Weimer ‘05 empirical model given steady state driving, focusing here on ionospheric quantities of interest. © Imperial College 2022

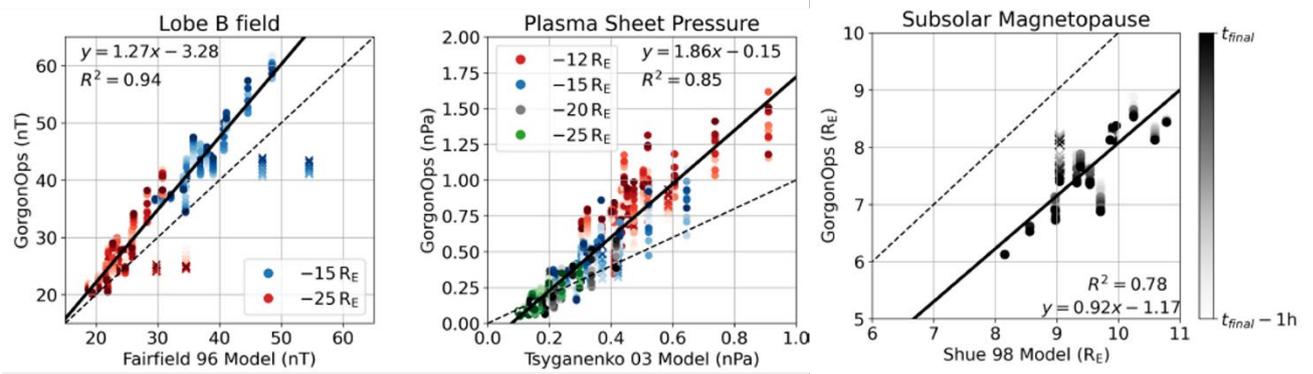


Figure 27: GorgonOps benchmarking relative to various empirical models for magnetospheric quantities of interest. © Imperial College 2022

Benchmarking for operational forecasts present an entirely different challenge, where the driving real-time solar wind data impacts the effectiveness of the forecast. Additionally, each separate forecast horizon should be judged individually, with the total performance being assessed across a period. An example of such validation is seen in Figure 28.

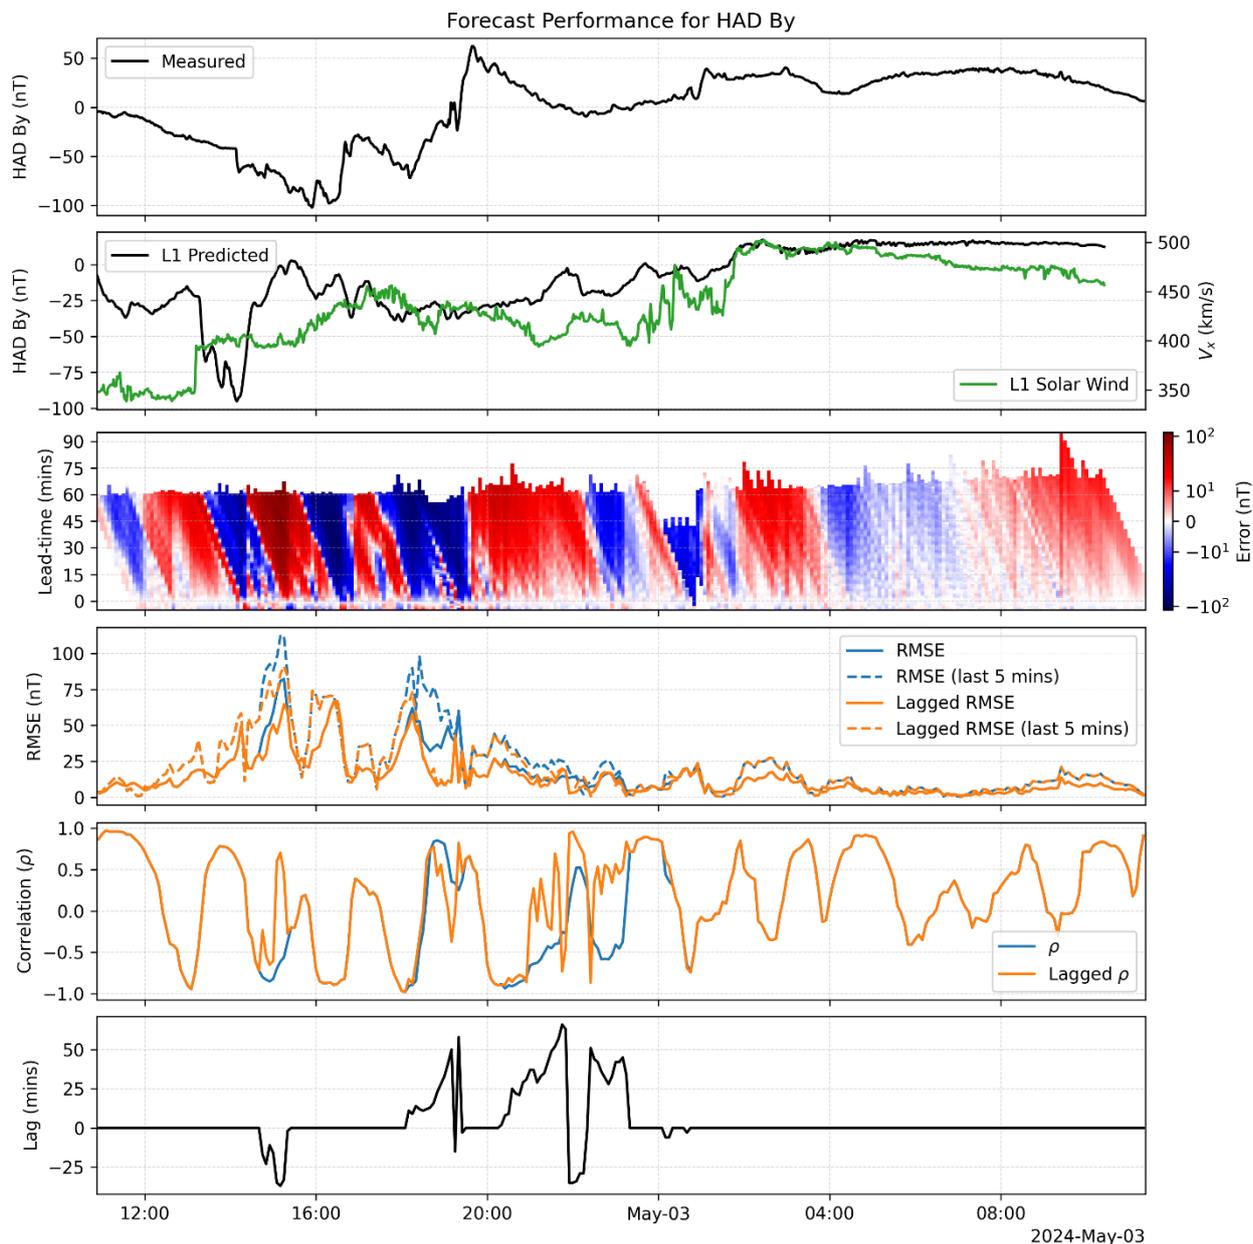


Figure 28: Forecast performance for GorgonOps during a G3 geomagnetic storm as measured at Hartland, with the By component of the magnetic field shown here. (Top panel) Measured real-time magnetometer data from the given station. (2nd panel) The full GorgonOps estimated profile, i.e., as if continuously run from L1 data ingested. The solar wind speed which informs the time horizon for a given simulation segment is added to this panel. Both profiles can be thought of as the raw profiles in the past. Note that the measured real-time field is used to estimate a baseline (at $t=0$) and 5 minutes prior to the forecast time is included in visualisation. (3rd panel) Error associated with each forecast (produced every 5 mins) for the given lead time determined by the solar wind speed. Heat map gives the error between predicted and measured for each subset. Here we can see the structures propagate, but do bear in mind that this shift is naive, and we are aiming to get the last 5 minutes as close as we can (but it is useful to see the profile for the forecast window). (4th panel) RMSE measures for the forecast window, last 5 minutes of the forecast window, and a shifted forecast that assumes the propagation using wind speed is not accurate (see 6th panel). (5th panel) Correlation coefficients for the forecast window and shifted forecast. (6th panel) Assuming the propagation using solar wind speed is not great, the lag estimated using cross correlation is shown here. This is used to shift the forecast in the above panels. © Imperial College 2024

5.3.2.1 GORGONOPS – CONFIGURATION

A brief description of the GorgonOps configuration is given in the table below. To note, all these parameters allow for robust operational running of the simulation of around 200 processor cores. These can be tweaked via a control file depending on user needs.

Operational Gorgon Setup	
Domain	$X = [-30, 130] R_E, Y = [-60, 60] R_E, Z = [-60, 60] R_E$ Inner boundary at $3.5 R_E$ $0.25 R_E$ inner boundary resolution with stretched grid implementation
Field	Dipole strength and latitude taken from IGRF data covering 1900-2025
Ionosphere	Empirical solar Extreme UV conductances, with MHD derived auroral, and background polar cap $\Sigma_P = 10$ mho, $\Sigma_H = 5$ mho
Solar Wind	Measured (e.g. RTSW/OMNI) or modelled (e.g. EUHFORIA) L1 data in GSM coordinates with Earth's diurnal rotation imposed

Table 3: GorgonOps configuration.

5.3.2.2 GORGONOPS – CONTINUOUS RUNNING

GorgonOps has been configured to optimise initialising and restarting simulations. Firstly, a library of preinitialised simulation states with a magnetosphere already formed are kept on-hand. These are specific to the time of year and the associated dipole tilt. For completeness, the initialization is made up of a 2-hour period where the solar wind driving fluctuates from $B_z = [2, -3, 2]$ nT. The last timestep of this initialization is taken as the simulation seed, and the new input solar wind data is ramped up from the initialization solar wind conditions over a period of 15 minutes. As noted elsewhere, there is a risk of input solar wind data not being available.

To overcome this and allow for continuous running of the simulation, various cold or warm restarts are automatically triggered (depending on the ingestion of real-time solar wind data). The last valid input timestep is tracked and used to assess the level of data outage or corruption. In each case, a restart makes use of a ramp up from the previous valid time instance to avoid spurious shocks. Briefly the approach to dealing with data gaps is defined as:

- Data gap < 15 mins: Continue waiting for new data (may reduce lead time)
- Data gap < 45 mins: Continue with last valid time instance, i.e., constant driving
- Data gap < 2 hours: Warm restart from last valid time instance, i.e., clean-up & wait
- Data gap > 2 hours: Cold restart from initialization, i.e., clean-up

5.3.2.3 GORGONOPS – GROUND GEOMAGNETIC FIELD ESTIMATION

As noted in the various containers available, during the SAGE project the capability to estimate the ground geomagnetic field using the Complex Image Method (CIM) for a simulation was developed. This and other functionality is ported within the Gorgon-Tools post-processing suite that can be accessed in future development. That said, validation of the method, in comparison to a highly efficient development of line-of-sight masking Biot-Savart integration over the simulation domain showed some shortcomings. Amongst these, is the lack of field-aligned current (FAC) contributions and the general computational overhead required. Furthermore, the ground induced current contribution introduced by CIM can be induced through an equivalent current approximation in the Biot-Savart integration approach. Figure 29 shows these considerations and Table 4 summarises the differences.

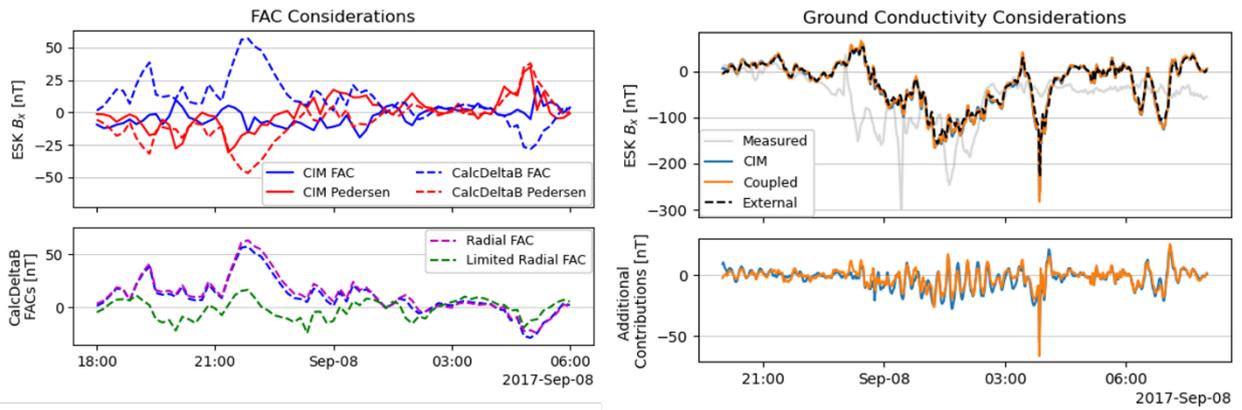


Figure 29: GorgonOps ground geomagnetic field estimation considerations. © Imperial College 2024

	CalcDeltaB	CIM
Contributions	Magnetosphere, FACs, ionosphere	Ionosphere, nett FACs on grid
FACs	Dipolar, line-of-sight	Vertical
Extent	Full simulation domain, spherical coords, line-of-sight	Stereographic projection with 1000km ² grid
Induced Current	Simplified CIM coupling using equivalent current	Explicit inclusion of ionospheric and FAC contributions

Table 4: Comparison of line-of-sight adapted CalcDeltaB (Biot-Savart integration) and Complex Image Method (CIM) methods.

5.4 NOTE ON USING REAL TIME SOLAR WIND DATA

Smith et al (2022) undertook a thorough study of the availability and quality of L1 data from the scientific releases on the OMNI system compared to the available data in real time (within 1-5 minutes of measurement). Their research showed large time gaps in the availability of some parameters, particularly plasma density. They also showed post-collection corrections occur which can change the magnitude or orientation of the data. Figure 30 shows an example of the change in the operational (NRT) versus science data from ACE and DISCOVER for the plasma density and velocity parameters. A perfect correlation would follow the 1:1 ratio dotted line within $\pm 10\%$ orange dotted lines.

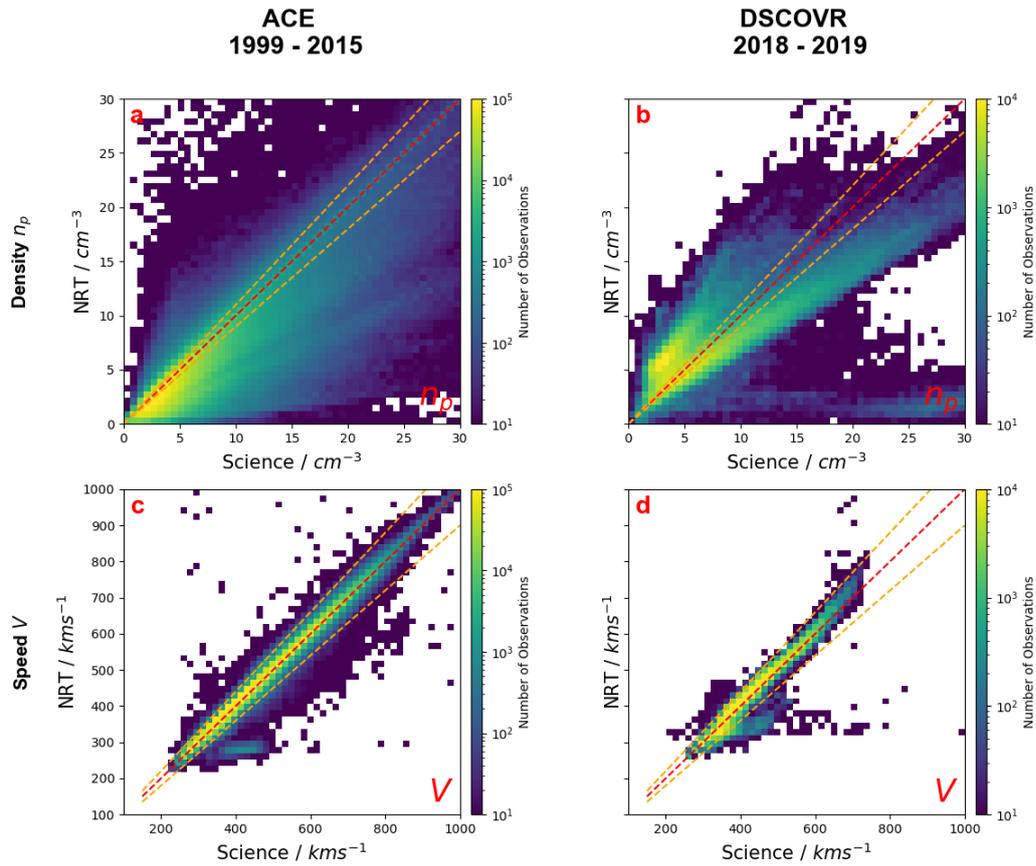


Figure 30: Comparison of the plasma moments in the Near-Real-Time (NRT) data with the science quality data. Left column (a, c) shows Advanced Composition Explorer and between 1999 and 2015 (inclusive), right column (b, d) shows Deep Space Climate Observatory between 2018 and mid-2019. The plasma moments shown are the plasma density (a, and b) and bulk velocity (c, and d). The diagonal dashed line indicates where the data returned are equivalent, while the orange dashed lines indicate the region where the NRT values are within $\pm 10\%$ of the science data. © MSSL/UCL 2022

They also showed that small short gaps are common in the operational data, and that interpolation of some of the data would vastly improve the up-time of the models that require long continuous time-series to operate. Figure 31 shows how interpolating short gaps of up to 15 minutes helps extend the availability of longer periods of data windows.

For example, the SPIDER model requires at least 120 minutes of continuous data to operate. If there is no interpolation applied, then using ACE will only provide forecasts less than 1%. A small amount of interpolation (five minutes) raises this to 75%.

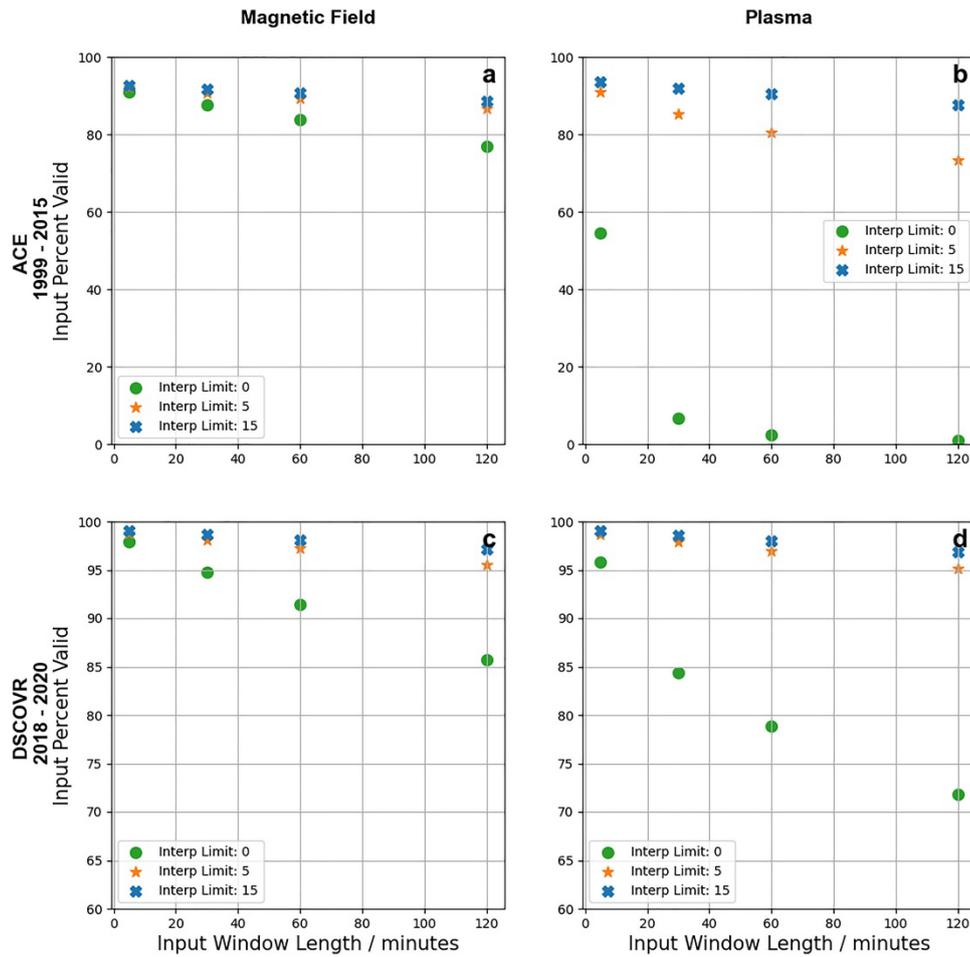


Figure 31: Fraction of data windows that are continuous, without data gaps for the Advanced Composition Explorer (top: a and b) and Deep Space Climate Observatory (bottom: c, d) Near-Real-Time data. The results for the magnetic field (left: a and c) and plasma data (right: b and d) are shown. The fraction of complete data windows are provided as a function of input window length required to be continuous. Three different interpolation schemes are presented: no interpolation (green circles), interpolation of gaps 5 min or shorter (orange stars) and interpolation of gaps 15 min or shorter (blue crosses). © MSSL/UCL 2022

References

- Beggan, C. D., Richardson, G. S., Baillie, O., Hübert, J., & Thomson, A. W. (2021). Geoelectric field measurement, modelling and validation during geomagnetic storms in the UK. *Journal of Space Weather and Space Climate*, 11, 37.
- Boteler, D. H. (2013). A new versatile method for modelling geomagnetic induction in pipelines. *Geophysical Journal International*, 193(1), 98-109.
- Boteler, D. H. (2021). Modeling geomagnetic interference on railway signaling track circuits. *Space Weather*, 19(1), e2020SW002609.
- Eroshenko, E. A., Belov, A. V., Boteler, D., Gaidash, S. P., Lobkov, S. L., Pirjola, R., & Trichtchenko, L. (2010). Effects of strong geomagnetic storms on Northern railways in Russia. *Advances in Space Research*, 46(9), 1102-1110.
- Florczak, E., Beggan, C. D., & Whaler, K. A. (2023). The predictive power of magnetospheric models for estimating ground magnetic field variation in the United Kingdom. *Frontiers in Astronomy and Space Sciences*, 10, 62.
- Forsyth, C., Rae, I. J., Coxon, J. C., Freeman, M. P., Jackman, C. M., Gjerloev, J., & Fazakerley, A. N. (2015). A new technique for determining Substorm Onsets and Phases from Indices of the Electrojet (SOPHIE). *Journal of Geophysical Research: Space Physics*, 120(12), 10-592.
- Gordeev, E., V. Sergeev, I. Honkonen, M. Kuznetsova, L. Rastätter, M. Palmroth, P. Janhunen, G. Tóth, J. Lyon, and M. Wiltberger (2015), Assessing the performance of community-available global MHD models using key system parameters and empirical relationships, *Space Weather*, 13, 868–884, doi:10.1002/2015SW001307.
- 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.
- Hübert, J., Beggan, C. D., Richardson, N. Gomez-Perez, A. Collins, & Thomson, A. W. P. (2024). Validating a UK Geomagnetically Induced Current Model using Differential Magnetometer Measurements, *Space Weather*, 22(2), e2023SW003769
- 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.
- Patterson, C. J., Wild, J. A., & Boteler, D. H. (2023). Modeling the impact of geomagnetically induced currents on electrified railway signaling systems in the United Kingdom. *Space Weather*, 21(3), e2022SW003385.
- Pulkkinen, A., et al. (2013), Community-wide validation of geospace model ground magnetic field perturbation predictions to support model transition to operations, *Space Weather*, 11, 369-385, doi:10.1002/swe.20056.

Smith, A.W., C. Forsyth, I.J. Rae, T.M. Garton, C.M. Jackman, M. Bakrania, R.M. Shore, G.S. Richardson, C.D. Beggan, M.J. Heyns, J.P. Eastwood, A.W.P. Thomson & J.M. Johnson. (2022). On the considerations of using near real time data for space weather hazard forecasting. *Space Weather*, 20(7), e2022SW003098.

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, 1605–1617. 2019SW002281.

Smith, A. W., Forsyth, C., Rae, I. J., Garton, T. M., Bloch, T., Jackman, C. M., & 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, e2021SW002788