

UNIVERSITY OF LEEDS

# Flows and Forecasts: Investigating the Geomagnetic Field using Machine Learning

Naomi Mara Shakespeare-Rees

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds

Faculty of Environment

School of Earth and Environment

August 2026

# Intellectual Property

The candidate confirms that the work submitted is her own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

© 2026 The University of Leeds, Naomi Mara Shakespeare-Rees

Signed



# Declaration

This research has been carried out by a team which has included Philip Livermore, Christopher Davies, Hannah Rogers, Ciarán Beggan, William Brown, and Christopher Finlay. I declare that the thesis has been composed by myself and that the work has not be submitted for any other degree or professional qualification. I confirm that the work submitted is my own except where work has formed part of jointly-authored publications.

The work presented in Chapter 4 was previously published in Physics of Earth and Planetary Interiors (2025) as 'Local flow estimation at the top of the Earth's core using Physics Informed Neural Networks' by '**Naomi Shakespeare-Rees**, Philip W. Livermore, Christopher J. Davies, Hannah F. Rogers, William J. Brown, Ciarán D. Beggan, and Christopher C. Finlay', as part of the special issue 'Terrestrial Magnetism: Insights from ten years of the Swarm satellite constellation'. This study was conceived by all the authors. I wrote a new codebase for this study, which is now published as a python package on github, and all the co-authors helped with analysis and discussion of the results.

The work presented in Chapter 5 was previously published in Earth, Planets and Space as 'Forecasting Secular Variation using Physics-Informed Neural Networks for IGRF-14' by '**Naomi Shakespeare-Rees**, Philip W. Livermore, Christopher J. Davies, Hannah F. Rogers, William J. Brown, Ciarán D. Beggan, and Christopher C. Finlay', as part of the special issue 'International Geomagnetic Reference Field - The Fourteenth Generation'. This study was conceived by all the authors. I adapted code that I wrote for Chapter 4, and the results from this study were used as a candidate SV forecast for the 14th International Geomagnetic Reference Field.

# Acknowledgements

This thesis would not be possible without the enormous help of the people around me, who have supported me, guided me, and/or dragged me to the pub or up a mountain.

Firstly, I would like to thank my small army of supervisors, Phil Livermore, Chris Davies, Hannah Rogers, Will Brown, Ciarán Beggan and Chris Finlay. Your wisdom, knowledge and patience has meant so much to me over the last four years, and I cannot thank you enough. I'd also like to thank others in the British Geological Survey's geomagnetism team, particularly Gemma, Laura and Ciara. You made my placement at BGS such an insightful and fun experience, from which I learnt so much. I'd also like to thank my academic sibling Fred, for the many discussions, helpful or otherwise, that we have had over the years. Long may the mutual bullying of Hannah continue. I'd also like to thank Ben Craven, without whom I would not have a working computer to write this thesis.

I am also profoundly grateful for the friends I have who have cheered me on during this process, both those made during the PhD and before. Thomas and Zoe, you have been an unwavering source of laughter, advice, and frankly *rude* comments since our days sat at the back of physics lectures at Glasgow, and I'm so happy we are now all at the end of our PhD journey. Izzi, you have always had my back since day one, and I am so grateful for our years of friendship. Thank you to Zak, Jared, Neil and Freya for the weekly hilarity, one day I may learn the rules to the game we have played for almost a decade. Lily, Helen, Rosie, Laura, and Mia, our wine nights have carried me through this PhD, and I hope they continue even as we move on from Leeds. Thank you to Ella, Will, and Stijn, for the constant laughs, and to Imogen, Joe and Bob, for always being up for outdoor adventures. Thank you also to everyone at Warrior Woman BJJ and Kickboxing, particularly my coach Jackie, for reminding me that there is a life outside academia.

Finally, I would like to extend my deepest gratitude to my family. To Mum and Mark, Dad and Jenny, thank you for your generous support and encouragement. Mum's scathing review of my first paper, "It seems like blobs to me, but I'm sure it seems like more than blobs to you" still comes to mind whenever I see any image of the geomagnetic field, so cheers for that Mum. Thank you to my sister Martha, for calling me a nerd constantly, and to my cats, for screaming at me when I stay up too late. Last, but certainly not least, thank you to my fiancée Cazza. Without your love, wisdom, humour, and encouragement, none of this would have been possible, and I can't wait to find out what our future brings.

# Abstract

The Earth's magnetic field has wide ranging impacts on navigation, satellite operations and protection from space weather, making it important to understand its generation and forecast its evolution. The dominant surface field is generated by electrically conducting fluid motion in the outer core through the geodynamo. Over short timescales, these field changes can be approximated as advection-dominated, relating geomagnetic secular variation to core-surface flow. However, this inversion is underdetermined and requires physical constraints. Typical approaches to core flow inversions are to use global spherical harmonic representations, which can couple separated regions and smooth localised features. Regional inversions avoid this coupling, allowing a test of whether inferred structures are supported by local magnetic-field information. Geomagnetic behaviour may also be influenced by heterogeneous heat flux at the core–mantle boundary (CMB). Geodynamo simulations suggest that this forcing affects magnetic-field morphology although its structure and amplitude remain poorly constrained.

This thesis describes the development of novel machine-learning approaches for investigating core-surface flow and CMB heat-flux heterogeneity. First, Physics-Informed Neural Networks (PINNs) are used to infer regional flows that fit local secular variation while satisfying a physical constraint, which are used to produce a mosaic of independent regional flows that recovers the planetary gyre. The PINN framework is then used to construct mesh-free regional and global flow models for geomagnetic forecasting. Hindcasts for 2020–2025 show that the regional approach produces lower root-mean-square misfits relative to CHAOS-7.18 than the other IGRF-13 candidate models. The method is subsequently used to produce a candidate secular-variation forecast for IGRF-14 over 2025–2030. Finally, a convolutional autoencoder classifier is developed to infer CMB heat-flux heterogeneity from the radial magnetic field. The results show that this heterogeneity leaves classifiable information in the field, but further development is required before the method can reliably constrain the Earth's CMB heat-flux structure.

# Contents

**Intellectual Property** **ii**

**Declaration** **iii**

**Acknowledgements** **iv**

**Abstract** **vi**

**Motivation and Structure** **xvii**

References . . . . . xix

**1 Geomagnetism and Outer Core Dynamics** **1**

1.1 Measurements of the Geomagnetic Field . . . . . 1

1.1.1 Satellite Data . . . . . 4

1.1.2 Reconstructions of the Ancient Field . . . . . 6

1.2 Spherical Harmonic Representation of the Geomagnetic Field . . . . . 8

1.3 Earth's Structure . . . . . 10

1.4 Sources of the Geomagnetic Field . . . . . 11

1.4.1 The Core Field . . . . . 14

1.4.2 The Lithospheric Field . . . . . 19

1.4.3 The External Field . . . . . 22

1.5 Modelling and Predicting the Geomagnetic Field . . . . . 24

1.5.1 The International Geomagnetic Reference Field . . . . . 25

1.5.2 The CHAOS Field Models . . . . . 26

1.6 Relating Secular Variation to Core Flow . . . . . 28

1.7 Poloidal–Toroidal Decomposition . . . . . 30

1.8 Flow Non-Uniqueness . . . 31
1.8.1 Tangentially Geostrophic Flow . . . 33
1.8.2 Other Flow Constraints . . . 34
1.9 Comparison of Regional and Global Flow . . . 35
1.10 Numerical Simulations . . . 37
1.11 Core–Mantle Boundary Coupling . . . 40
1.11.1 Large Low-Velocity Provinces (LLVPs) . . . 40
1.11.2 Outermost-Core Stratification . . . 41
1.11.3 Heat Flux Heterogeneity . . . 42
References . . . 44

**2 Machine Learning** **57**
2.1 Model Training and Evaluation . . . 59
2.1.1 Loss Functions . . . 59
2.1.2 Optimisation . . . 61
2.1.3 Hyperparameters . . . 62
2.1.4 Data Scaling . . . 64
2.1.5 Model Evaluation . . . 64
2.1.6 Generalisation . . . 66
2.2 Machine-Learning Architectures . . . 68
2.2.1 Neural Networks . . . 68
2.2.2 Physics-Informed Neural Networks . . . 71
2.2.3 Convolutional Layers . . . 74
2.2.4 Autoencoder Classifiers . . . 76
2.3 Explainable Artificial Intelligence . . . 78
References . . . 81

**3 Research Questions** **86**
References . . . 87

**4 Local Flow estimation at the top of the Earth's core using Physics Informed Neural Networks** **89**
4.1 Introduction . . . 91

4.2 Method . . . 94
4.2.1 Mathematical Framework . . . 94
4.2.2 Neural Networks . . . 96
4.2.3 Training Dataset: CHAOS-8.1 . . . 98
4.2.4 Parameter Choices . . . 99
4.2.5 Synthetic Tests . . . 103
4.3 Results . . . 106
4.3.1 Atlantic . . . 106
4.3.2 Global Patchwork . . . 108
4.3.3 Equatorial Flows . . . 110
4.3.4 High Latitude Jet near the North Pole . . . 114
4.4 Discussion . . . 116
4.5 Conclusion . . . 117
References . . . 118

**5 Forecasting Secular Variation using Physics-Informed Neural Networks 123**
5.1 Introduction . . . 125
5.2 Core Flow Inversion using PINNs . . . 128
5.2.1 Training Data . . . 131
5.3 Advection Methodology . . . 134
5.4 Hindcasting and comparison to IGRF-13 . . . 135
5.4.1 Regional and Global Models . . . 135
5.4.2 Comparison to IGRF-13 . . . 140
5.5 Candidate SV Model . . . 143
5.6 Discussion and Conclusions . . . 146
5.7 Summary . . . 148
5.A Updated Gauss Coefficients for Leeds' Candidate SV Model . . . 149
5.B Comparison between IGRF-13 Candidates and Kalmag . . . 150
References . . . 151

**6 Using autoencoders to classify geodynamo heat flux heterogeneity from snapshots of the magnetic field 156**
6.1 Introduction . . . 158

6.2 Methodology . . . 162
6.3 Classification using simulations in Mound and Davies (2023) . . . 166
6.3.1 Training Data . . . 166
6.3.2 Hyperparameter Tuning . . . 169
6.3.3 Sensitivity to regions of the CMB . . . 175
6.3.4 Saliency Maps . . . 178
6.4 Classification using a wider set of simulations . . . 184
6.5 Classification using models of the Earth . . . 192
6.6 Discussion and Conclusions . . . 198
6.7 Summary . . . 201
References . . . 215

**7 Conclusions** **219**
7.1 Can regional core-flow inversions recover the large-scale structures found in global flow models? . . . 220
7.2 Do localised flow models contain information helpful for forecasting the Secular Variation of the Earth's magnetic field? . . . 223
7.3 Can heat flux heterogeneity be inferred from the radial magnetic field at the Core Mantle Boundary? . . . 226
7.4 Concluding Remarks and Further Work . . . 228
References . . . 230

**A Appendix Chapter** **233**
A.1 Toroidal Flows . . . 233
A.2 TG Misfit . . . 235

# List of Figures


1.1 Components of the geomagnetic field vector (Hulot et al., 2015). . . . . 2

1.2 Map of INTERMAGNET observatories, (INTERMAGNET) . . . . . 3

1.3 Comparison of orbit trajectories (Gao et al., 2025). . . . . 5

1.4 Schematic of a satellite passing through a Geomagnetic Virtual Observatory (GVO) footprint (Hammer et al., 2021) . . . . . 6

1.5 Schematic picture of the Earth, not to scale. . . . . 11

1.6 Sketch of the various sources contributing to the near-Earth magnetic field (Olsen et al., 2010b). . . . . 12

1.7 The Grand Spectrum of the geomagnetic field (Constable and Constable, 2023). . 13

1.8 Radial component of the main field, shown using the CHAOS-8 model (Kloss et al., 2026). . . . . 16

1.9 Radial component of the Secular Variation, shown using the CHAOS-8 model (Kloss et al., 2026). . . . . 17

1.10 Time–longitude plot of the radial component of secular variation at the CMB along the equator. . . . . 18

1.11 The first version of the World Digital Magnetic Anomaly Map (WDMAM) (Thébault et al., 2010). . . . . 20

1.12 Geomagnetic power spectra for the NGDC-720 model (Maus, 2008). . . . . 21

1.13 Schematic of the Earth's magnetosphere and the principal external current systems (Kelley and Heelis, 1989). . . . . 23

1.14 Time–longitude plots of the radial secular variation, for CHAOS-8.1 and CHAOS-7.18, (Kloss et al., 2026). . . . . 27

1.15 Tangentially geostrophic flow (a) and toroidal flow (b) derived from magnetic satellite data for 1999–2004. Adapted from Holme and Olsen (2006). . . . . 32

1.16 Schematic of structures in the core–mantle boundary region, (Hernlund and McNamara, 2015). . . . . 40

2.1 Schematic loss function showing a local minima and global minima (Goodfellow et al., 2016). . . . . 62
2.2 Schematic showing the effect of the learning rate on loss minimisation. . . . . 63
2.3 Example confusion matrix for a binary classification problem. . . . . 65
2.4 Examples of underfitting, an appropriate model capacity and overfitting for a quadratic dataset (Goodfellow et al., 2016). . . . . 66
2.5 Example schematic of a neural network, showing the input layer, hidden layers and output layer. . . . . 69
2.6 Example of two-dimensional cross-correlation without kernel flipping (Goodfellow et al., 2016). . . . . 75
2.7 Schematic of an autoencoder (Bergmann and Stryker, 2026). . . . . 76
2.8 Example Grad-CAM visualisations for a visual-question-answering model (Selvaraju et al., 2017). . . . . 80

4.1 Schematic of the PINN used in Chapter 4. . . . . 96
4.2 Example of a latitude-longitude box used for inversion. . . . . 99
4.3 Trade-off curves for size of network tests. . . . . 101
4.4 Recovered flows and SV Residuals at the CMB for the network sizes that produce the least complex and most complex flows, in the area under the Atlantic. . . . . 102
4.5 Recovered flows and SV Residuals at the CMB for the network sizes that produce the least complex and most complex flows, in the area under the South China Sea. 102
4.6 Example of the PINN recovery of the flow, based on a synthetic test using rigid body rotation. . . . . 104
4.7 Loss plotted against iteration. . . . . 107
4.8 (A) Recovered Flow under the northern Atlantic, at 35°N to 5°N latitude, showing westward flow. (B) Recovered SV (C) Residuals between the SV from CHAOS-8.1 and the SV shown in (B). . . . . 107
4.9 Global patchwork of regional flows, Recovered SV, and Residuals. . . . . 109
4.10 Time-Longitude plot at the equator, spanning 2000.0 to 2024.0. . . . . 111

4.11 Changes in the horizontal flow under Indonesia during the period 2010.0 to 2024.0, at 10° above and below the equator. . . . . 112
4.12 Residual between input SV and output SV for 2024, with contours of $B_r/\cos\theta$ superimposed. . . . . 113
4.13 Changes under the Eastern Pacific during the period of 2020.0 to 2024.0, at 10° above and below the equator. . . . . 113
4.14 Residual between input SV and output SV for 2024, with contours of $B_r/\cos\theta$ superimposed. . . . . 114
4.15 High Latitude jet in 2017.0 and 2024.0. . . . . 115
4.16 Maximum azimuthal flow over time. . . . . 115

5.1 Schematic of the PINN used in Chapter 5. . . . . 129
5.2 SV from CHAOS-7.18 at epoch 2020.0 (Finlay et al. (2020)) at the CMB. . . . . 133
5.3 Example of a box used for inversion. . . . . 133
5.4 Regional and Global flow models. . . . . 136
5.5 Average SV over the period 2020–2025, from CHAOS-7.18 model, 'Global' flow model, 'Regional' flow model, and Absolute difference between hindcasts. . . . . 137
5.6 Mausberger-Lowes spectrum of the SV for the global and regional hindcasts. . . 139
5.7 Mausberger-Lowes spectrum for our global and regional hindcasts, as well as the other IGRF-13 candidate models. . . . . 141
5.8 RMS between IGRF-13 candidates and CHAOS-7.18 derived $B_r$ values over time. 142
5.9 Recovered SV, Residual between CHAOS-7.18 SV and the Recovered SV, and Recovered Flows from the PINN . . . . . 144
5.10 Predicted SV (Top: $\partial B_r/\partial t$, Middle: $\partial B_\theta/\partial t$, Bottom: $\partial B_\phi/\partial t$) for epochs 2025.0-2030.0 at the core-mantle boundary. Continents shown for reference. . . . 145
5.11 RMS between IGRF-13 candidates and Kalmag derived $B_r$ values over time. . . 150

6.1 Schematic of the convolutional autoencoder classifier used in Chapter 6. . . . . 165
6.2 Class-averaged radial magnetic field in the equatorial region, truncated at spherical harmonic degree 12. . . . . 167
6.3 Class-averaged radial magnetic field at the poles, truncated at spherical harmonic degree 12. . . . . 168

6.4 Classification error (%) for the validation dataset against spherical harmonic degree, for the three-class approach (top) and the two-class approach (bottom). . 171
6.5 Classification percentage error against compression ratio. The mean classification error is shown in blue, and the individual results from each random initialisation are shown in black. . . . . 173
6.6 Classification Error when different regions were used for training data. . . . . 177
6.7 Example of an individual saliency map, with the radial magnetic field plotted on the same grid but using a different colour map. . . . . 179
6.8 Aggregate saliency maps for the three-class classification approach, shown in a confusion-matrix layout. . . . . 180
6.9 Map showing the heat flux pattern of the LLVPs applied to the training simulations. . . . . 180
6.10 Aggregate saliency maps for the two-class classification approach, shown in a confusion-matrix layout. . . . . 182
6.11 Examples of individual saliency maps for a $Q^* \neq 0$ 2-class case, with the radial magnetic field plotted alongside . . . . . 183
6.12 Class-averaged radial magnetic field in the equatorial region, truncated at spherical harmonic degree 12. . . . . 188
6.13 Class-averaged radial magnetic field in the polar region, truncated at spherical harmonic degree 12. . . . . 188
6.14 Aggregate saliency maps using the full set of simulations in Table 6.2, shown in a confusion-matrix layout and aggregate saliency maps using heterogenous simulation where $Q^* > 1$ only. . . . . 191
6.15 Mean normalised Lowes–Mauersberger power spectra for the homogeneous and heterogeneous simulations used for training, and for the CHAOS-8 and gufm1 Earth magnetic-field models. . . . . 193
6.16 Aggregate saliency map for the CHAOS-8.5 classifications, with the boundary of the South Atlantic Anomaly, defined as the contour around the area below 30,000nT at the Earth's surface, shown in green. . . . . 194
6.17 [Top] Aggregate saliency map for snapshots classified as homogenous, [Bottom] Aggregate saliency map for snapshots classified as heterogenous. . . . . 197

7.1 Mauersberger–Lowes spectra of the secular-variation forecast candidates for IGRF-14. The mean and median of the candidates are shown by the solid green and purple lines, respectively, while the updated Leeds candidate is shown by the solid black line. Figure created by adapting code from the IGRF-14 evaluation repository (IAGA V-MOD). . . . . 224

7.2 Difference between each IGRF-14 candidate radial secular-variation forecast and the median of all candidate forecasts. Figure created by adapting code from the IGRF-14 evaluation repository (IAGA V-MOD), plotted at the Earth's surface. . 225

A.1 (Top) Global patchwork of toroidal regional flows in January 2020. The thickness of the arrows indicate flow magnitude, and the arrows indicate direction. (Middle) SV calculated using the recovered flows in the top panel. (Bottom) Residuals between the SV from CHAOS-8.1 and the SV shown in the middle panel. The box boundaries (after the 5◦ border removal) are marked by grey dashed lines. . 234

A.2 TG misfit for the global mosiac shown in Figure 4.9, in squared units as the misfit is quantified by the MSE. . . . . 235

# List of Tables

1.1 Summary of low Earth orbit satellites with scientific magnetometer payloads. . 5

1.2 Summary of dimensionless control parameters used in Equations (1.35)-(1.38). . 39

2.1 Summary of commonly used neural-network activation functions. . . . . 71

4.1 Synthetic flows used for testing. . . . . 105

5.1 Updated Gauss coefficients for the University of Leeds IGRF Candidate model. . 149

6.1 Summary of the simulations used for training (Mound and Davies, 2023). . . . . 167

6.2 Summary of the simulations used for training data. . . . . 185

6.3 Results for wider simulation training data. . . . . 189

6.4 Classification results for the individual gufm1 snapshots and the time-averaged field using classifiers initialised with five different random seeds. . . . . 195

6.4 Classification results for the individual gufm1 snapshots and the Time-Averaged Field (TAF) using classifiers initialised with five different random seeds. . . . . 196

6.5 Reference and forcing metadata for the simulations listed in Table 6.2. . . . . 202

# Motivation and Structure

The predominantly dipolar geomagnetic field extends from the Earth's deep interior to approximately 90,000 km above its surface (Cox and Brown, 2013). It comprises contributions from multiple internal and external sources, although up to 90% of the surface field is generated by fluid motion in the outer core through the geodynamo (Olsen et al., 2010). Because the Earth's interior cannot be observed directly, changes in the geomagnetic field provide one of the few means of investigating the dynamics of the outer core. The field varies over timescales from seconds to millennia, exhibiting secular variation, polar wandering and occasional polarity reversals (Jackson et al., 2000; Livermore et al., 2020; Panovska et al., 2019). However, important questions remain about the physical processes responsible for these variations, including the structure and evolution of core flow, the mechanisms driving rapid field changes, and the influence of the lowermost mantle on the geodynamo. Understanding these physical processes is therefore fundamental to constraining the dynamics and evolution of the Earth's deep interior.

Scientific machine learning is an emerging field that combines data-driven methods with physical knowledge, but its application to the study of the Earth's magnetic field remains relatively under-explored. This thesis investigates its potential for studying the geomagnetic field, focusing on the dynamics of the outer core and the influence of heat-flux heterogeneity at the core–mantle boundary. The structure of the thesis is as follows:

- **Chapter 1: Geomagnetism and Outer Core Dynamics**

  This chapter provides the geophysical background for the thesis. It introduces geomagnetic observations and their spherical harmonic representation, before discussing the different sources of the measured field. It then considers geomagnetic field modelling and prediction, the relationship between secular variation and core-surface flow, numerical geodynamo simulations, and thermal coupling across the core–mantle boundary.

- **Chapter 2: Machine Learning**

  This chapter introduces the machine-learning concepts used throughout the thesis. It describes model training, validation and evaluation, before introducing neural networks, physics-informed neural networks, convolutional architectures and autoencoders. It also considers how physical constraints can be incorporated into machine-learning models and how explainable AI methods can be used to investigate their predictions.

- **Chapter 3: Research Questions**

  This chapter explores the research questions that will be answered in this thesis.

- **Chapter 4: Local Flow estimation at the top of the Earth's core using Physics-Informed Neural Networks**

  This chapter describes the development of a regional core-flow inversion methodology using physics-informed neural networks. The method is validated using synthetic flows before being applied to CHAOS-8.1 (Kloss et al., 2026). Independently trained regional inversions are then combined to test whether large-scale and localised flow features identified by global models can be recovered using only regional magnetic-field and secular-variation information.

- **Chapter 5: Forecasting Secular Variation using Physics-Informed Neural Networks**

  This chapter uses the core-flow inversion methodology to forecast geomagnetic secular variation. Regional and global flow models are inferred from CHAOS-7.18 and used to advect the magnetic field under the assumption of steady flow (Finlay et al., 2020). Hindcasting experiments are used to compare their performance with the IGRF-13 candidate models, before the regional approach is used to produce the University of Leeds candidate forecast for IGRF-14.

- **Chapter 6: Using autoencoders to classify geodynamo heat flux heterogeneity from snapshots of the magnetic field**

  This chapter develops a convolutional autoencoder classifier to test whether CMB heat-flux heterogeneity can be inferred from radial magnetic-field snapshots. The method is first tested using controlled geodynamo simulations published in Mound and Davies (2023),

and then extends to a wider simulation set. Sensitivity tests and saliency maps are used to investigate the information contributing to the classifications, before the method is applied to the CHAOS-8.5 and gufm1 models of the Earth's magnetic field (Kloss et al., 2026; Jackson et al., 2000).

- **Chapter 7: Conclusions**

  This chapter brings together the findings of the three research chapters and discusses them in the context of the research questions introduced in Chapter 3. It discusses the novel contributions and limitations of the methodologies, before considering future developments needed to produce more robust, reproducible and physically interpretable applications of machine learning in geomagnetism.

# References

Cox, G., Brown, W., 2013. Rapid dynamics of the earth's core. Astronomy & Geophysics 54, 5.32–5.37. doi:10.1093/astrogeo/att167.

Finlay, C.C., Kloss, C., Olsen, N., Hammer, M.D., Tøffner-Clausen, L., Grayver, A., Kuvshinov, A., 2020. The chaos-7 geomagnetic field model and observed changes in the south atlantic anomaly. Earth, Planets and Space 72, 156. URL: https://doi.org/10.1186/s40623-020-01252-9, doi:10.1186/s40623-020-01252-9.

Jackson, A., Jonkers, A.R.T., Walker, M.R., 2000. Four centuries of geomagnetic secular variation from historical records. Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 358, 957–990. doi:10.1098/rsta.2000.0569.

Kloss, C., Finlay, C.C., Olsen, N., Tøffner-Clausen, L., Gillet, N., Grayver, A., 2026. The CHAOS-8 Geomagnetic Field Model. Earth, Planets and Space 78, 21. URL: https://doi.org/10.1186/s40623-025-02352-0, doi:10.1186/s40623-025-02352-0.

Livermore, P.W., Finlay, C., Bayliff, M., 2020. Recent north magnetic pole acceleration towards siberia caused by flux lobe elongation. Nature Geoscience 13, 387–391. doi:10.1038/s41561-020-0570-9.

Mound, J.E., Davies, 2023. Longitudinal structure of earth's magnetic field controlled by lower mantle heat flow. Nature Geoscience 16, 380–385. doi:10.1038/s41561-023-01148-9.

Olsen, N., Hulot, G., Sabaka, T., 2010. Sources of the Geomagnetic Field and the Modern Data That EnableTheir Investigation. Springer, Berlin, Heidelberg. p. 105–124. URL: https://doi.org/10.1007/978-3-642-01546-5_5, doi:10.1007/978-3-642-01546-5_5.

Panovska, S., Korte, M., Constable, C.G., 2019. One hundred thousand years of geomagnetic field evolution. Reviews of Geophysics 57, 1289–1337. doi:10.1029/2019RG000656.

# Chapter 1

# Geomagnetism and Outer Core Dynamics

Geomagnetism is the subject of the Earth's magnetic field, the processes that generate and modify it, and its evolution through time. This chapter reviews the pre-existing literature, beginning with geomagnetic observations before introducing the field's mathematical representation and principal sources. It then considers field modelling and prediction, core-flow inference, and numerical geodynamo simulations. The chapter concludes with core–mantle boundary structure and its influence on core dynamics and the geomagnetic field.

## 1.1 Measurements of the Geomagnetic Field

While knowledge of the existence of magnetism dates back to pre-history — with the earliest observations of natural magnets being attributed to the Greek philosopher Thales (624-546 BC) (Jonkers, 2003) — awareness of the Earth's magnetic field emerged at a slower pace. The first description of a magnetic needle compass for navigation was described by Shen Kuo in his Dream Pool essays, published in 1088 AD, with the first European description by Alexander Neckam in 1190 AD (Kreutz, 1984). By at least the 14th century, European sailors were well aware of declination, $D$, — the difference between geographic north and magnetic north, shown in Figure 1.1 — with Christopher Columbus making measurements of this during his sea voyage in 1492 AD. The publishing of *De Magnete* by William Gilbert in 1600 AD marked the first understanding that the magnetic field was of internal origin, deducing after experiments with a model Earth that 'the Earth globe itself is a great magnet' (Gilbert, 1600). Edmond Halley

published the first map of magnetic declination in 1701, in the form of an isogonic map with contour lines across the Atlantic Ocean connecting points where the compass had the same variation (Halley, 1701). By the mid-17th century it was recognised that the geomagnetic field varies in time, a phenomenon known as secular variation (Jonkers, 2003). Motivated by the effect that this variation had on navigation (Jackson et al., 2000), a network of ground-based geomagnetic observatories was established by 1837 to determine how the geomagnetic field changes spatially and temporally (Jackson and Finlay, 2015).

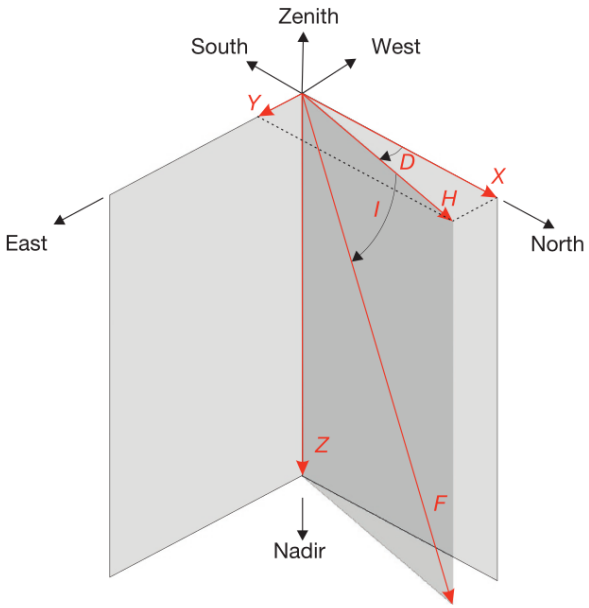



Figure 1.1: Components of the geomagnetic field vector, $\boldsymbol{B}$, in a right-handed coordinate system. The quantities $X$, $Y$ and $Z$ are the three vector field components, $H$ is the horizontal magnitude, $F$ is the total intensity, $D$ is the declination, and $I$ the inclination, the angle that the Earth's magnetic field lines make with the horizontal plane of the Earth's surface. Taken from Hulot et al. (2015).

Geomagnetic observatories typically use one or both of two types of magnetometer: scalar and vector (Gubbins and Herrero-Bervera, 2007). Scalar magnetometers measure only the total field intensity, whereas vector magnetometers measure both the magnitude and direction of the magnetic field. Measurement of the full three-dimensional field requires three orthogonal sensors, corresponding to the three field components marked as $X$, $Y$ and $Z$ in Figure 1.1. Magnetometers may also be described as absolute or relative instruments. Absolute instruments can be calibrated using known internal constants, while relative instruments require calibration against a known magnetic field (Rasson, 2007). To ensure consistent, high-quality, data the International Real-time Magnetic Observatory Network (INTERMAGNET) was founded in 1987. There are now 124 active INTERMAGNET observatories worldwide (as of June 2026), gener-

ating continuous time series of regular measurements of both the field intensity and direction. A map of the active INTERMAGNET observatories is shown in Figure 1.2. Unfortunately, the distribution of these observatories is highly uneven, being concentrated on land in the Northern Hemisphere, primarily due to geographical and political constraints. Each of these observatories must be in geomagnetically quiet areas, away from the influence of human activity, but at a location in which regular absolute and relative readings, as well as maintenance, can be carried out. The resulting gaps in the data coverage lead to uneven spatial coverage. The advent of satellite-based geomagnetic measurements therefore offers global sampling which complements, but does not replace, ground-based observatories.

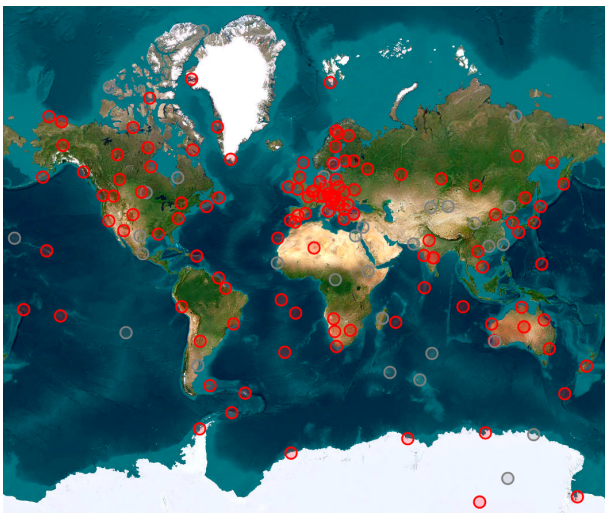


Figure 1.2: Map showing all the active INTERMAGNET observatories in red, with the closed observatories in grey. Taken from INTERMAGNET.

### 1.1.1 Satellite Data

Satellite measurements have been used since 1965 to complement ground-based observations, with continuous satellite records available since 1999. A summary of the satellites used to measure the geomagnetic field is provided in Table 1.1, while the orbital trajectories of the most utilised satellites for studying the geomagnetic field currently in operation are shown in Figure 1.3 (Gao et al., 2025). Magnetometers carried on board these satellites measure the geomagnetic field by revisiting approximately the same locations every few months. Repeat coverage helps to separate changes in the measured magnetic field caused by temporal evolution of the source from spatial variations encountered as the satellite moves. However, fully distinguishing these effects remains challenging because measurements at the same location are limited by the satellite's revisit rate. At present, three satellite missions with dedicated magnetometers for scientific gain remain operational. The oldest of these is ESA's Swarm mission, launched in 2013. Swarm comprises three identical low-Earth orbit satellites, Alpha, Bravo, and Charlie. Each satellite carries a Vector Field Magnetometer and an Absolute Scalar Magnetometer, in addition to instruments for navigation and electric field measurements (Friis-Christensen et al., 2006). The China Seismo-Electromagnetic Satellite (CSES), launched in 2018, consists of a single satellite in a polar, Sun-synchronous low-Earth orbit. It carries two fluxgate magnetometers, which provide vector measurements, and a Coupled Dark State Magnetometer (CDSM), which provides scalar measurements for calibration of the vector data (Pollinger et al., 2020). The Macau Scientific Satellite-1 (MSS-1), launched in 2023, consists of two satellites, MSS-1A and MSS-1B, in near-circular low-Earth orbits, meaning orbits relatively close to the Earth's surface, with an inclination of 41°, where inclination describes the angle between the orbital plane and the equatorial plane. MSS-1A carries a Coupled Dark State Magnetometer, together with two vector magnetometers, whereas MSS-1B carries instruments for energetic electron flux and solar X-ray observations (Jiang et al., 2024). In contrast to the polar orbits of Swarm and CSES, the lower-inclination orbit of MSS-1 is designed to provide rapid local-time coverage at low and mid-latitudes, particularly over the South Atlantic Anomaly (Jiang et al., 2024).

The movement of satellites with respect to the Earth's surface makes it difficult to construct time series at fixed locations, or to compare satellite measurements directly with ground-based observatory data. To address this, Mandea and Olsen (2006) proposed the Geomagnetic Virtual Observatory (GVO) method. In this approach, satellite measurements are collected within a

| Satellite Name | Operation Dates | Magnetometer Payload | Reference |
|---|---|---|---|
| OGO-2 | 10/1965 – 09/1967 | Scalar | Jackson and Vette (1975) |
| OGO-4 | 07/1967 – 01/1969 | Scalar | Jackson and Vette (1975) |
| OGO-6 | 06/1969 – 06/1971 | Scalar | Jackson and Vette (1975) |
| Magsat | 11/1979 – 05/1980 | Scalar and Vector | Langel et al. (1982) |
| Ørsted | 02/ 1999 – 01/2014 | Vector | Neubert et al. (2001) |
| CHAMP | 07/2000 – 09/2010 | Scalar and Vector | Reigber et al. (2002) |
| SAC-C | 01/2001 – 12/2004 | Scalar | Colomb and Varotto (2003) |
| Swarm | 11/2013 – Present | Scalar and Vector | Olsen et al. (2013) |
| CSES | 02/2018 - Present | Scalar and Vector | Zeren et al. (2025) |
| MSS-1A | 05/2023 - Present | Scalar and Vector | Li et al. (2025) |

Table 1.1: Summary of low Earth orbit satellites with scientific magnetometer payloads.

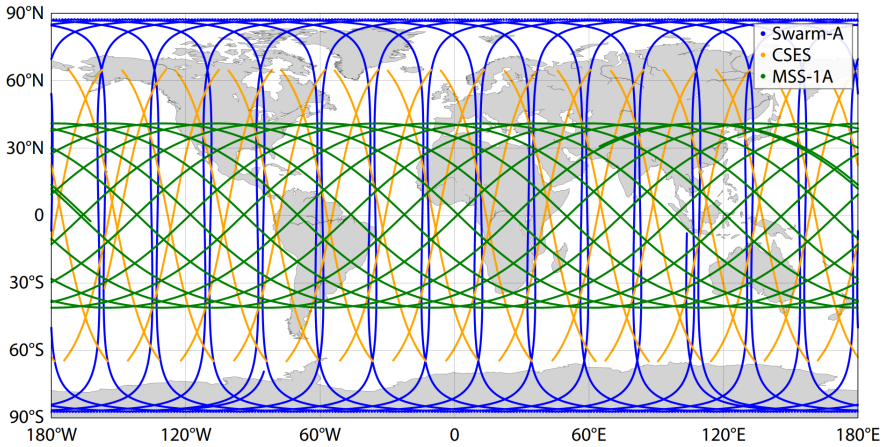



Figure 1.3: Comparison of orbit trajectories over one day for the Swarm A, CSES, and MSS-1A spacecraft. Taken from Gao et al. (2025).

specified radius of influence around a central point above the Earth's surface. A polynomial expansion of the magnetic potential, typically of quadratic or cubic order, is then used to estimate the magnetic field at that point. A schematic of this approach is shown in Figure 1.4. GVOs therefore provide a compromise between satellite and ground-based observations, combining the broader spatial coverage of satellite measurements with regularly sampled time series at fixed locations, although at a lower temporal resolution than ground-based observatories. The measurements may then be averaged over a given period or divided into time bins to produce regularly sampled time series at a fixed set of GVO locations. By reducing measurements collected along different satellite orbits to common locations, the method limits the influence of changes in orbital sampling between successive time steps, and allows for the combination of multiple satellite datasets (Brown et al., 2022). The resulting estimates provide an observatory-like representation of the field, but should not be treated as direct equivalents of ground-based observatory measurements because they are derived from spatially and temporally distributed

satellite data at altitude.

Despite these advantages, GVOs are not always the preferred data source. Because they are constructed from satellite measurements of the total geomagnetic field, they contain both internal and external field contributions. Separating these signals is a non-trivial problem and can introduce additional uncertainty. For the development and testing of new methodologies, researchers may therefore choose to use global field models instead, as these provide a more controlled estimate of the internal field. This can reduce the risk of noise or unmodelled external signals, particularly at polar latitudes, producing artefacts that could be difficult to isolate from methodological effects.

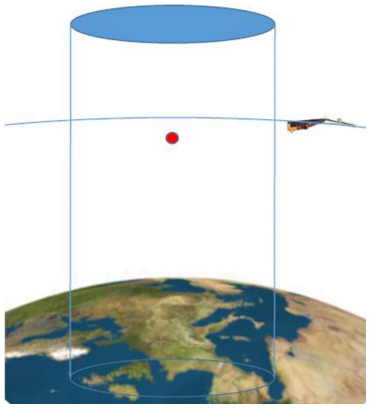


Figure 1.4: Schematic of a satellite passing through a Geomagnetic Virtual Observatory (GVO) footprint. Measurements collected along the satellite path, shown in blue, are reduced to an estimate of the magnetic field at the central GVO location, shown by the red dot. Taken from Hammer et al. (2021)

### 1.1.2 Reconstructions of the Ancient Field

Direct measurements of the geomagnetic field extend back only a few hundred years, representing a very small fraction of the lifetime of the field (which is at least 3.45 billion years (Usui et al., 2009)). Information about earlier field behaviour is therefore obtained using paleomagnetism, which studies the magnetic field recorded in rocks and sediments, and archaeomagnetism, which studies the magnetic field recorded in archaeological materials. These records arise from remanent magnetisation, whereby magnetic minerals preserve a signal of the ambient geomagnetic field acquired during their formation or during subsequent heating or alteration (Merrill et al., 1996; Tauxe, 2010; Kono, 2015). Paleomagnetic observations therefore provide the primary constraint on the evolution of the geomagnetic field over geological timescales (Panovska et al., 2019).

Paleomagnetic records are obtained from both volcanic and sedimentary materials, which provide complementary constraints on field behaviour. Volcanic sequences record the field at discrete times, while sediments provide more continuous records that can extend over hundreds of thousands of years, although with reduced temporal resolution and potential smoothing (Constable and Korte, 2015; Panovska et al., 2019). Paleomagnetic observations show that the geomagnetic field varies over a wide range of spatial and temporal scales. On centennial to millennial timescales, this behaviour is expressed as paleosecular variation, with changes in both direction and intensity reconstructed from volcanic, sedimentary, and archaeomagnetic data (Constable and Korte, 2015; Johnson and McFadden, 2015). When averaged over sufficiently long timescales, the field is well approximated by a geocentric axial dipole, although significant departures occur on shorter timescales and at regional scales (Johnson and McFadden, 2015). These observations provide constraints on core dynamics and the processes driving secular variation (Panovska et al., 2019).

In addition to paleosecular variation, paleomagnetic records document more extreme behaviour, including geomagnetic excursions and polarity reversals. Excursions correspond to short-lived departures from the existing polarity state, characterised by reduced dipole strength and increased field complexity, before recovery to the original polarity (Laj and Channell, 2015; Mason et al., 2024). Polarity reversals involve a complete change in the sign of the dipole component and have occurred hundreds of times throughout Earth's history, forming the geomagnetic polarity timescale (Jacobs, 1994; Merrill et al., 1996; Ziegler et al., 2011; Ogg, 2020). Reversals occur rapidly on geological timescales, typically over a few thousand years, and are associated with a significant reduction in dipole intensity (Valet, 2003). During these intervals the field becomes highly non-dipolar, with rapid changes in direction and intensity indicating large-scale reorganisation of the core field (Valet, 2003). Reconstructions over the last 100 kyr show substantial variability in dipole strength and document events such as the Laschamp excursion, highlighting the strongly time-dependent behaviour of the geodynamo (Panovska et al., 2019).

Paleomagnetic data sampling is spatially uneven and temporally incomplete. Volcanic records are episodic and geographically clustered, while sedimentary records may be smoothed and subject to age uncertainties (Constable and Korte, 2015). Nevertheless, paleomagnetism provides the only direct record of geomagnetic field behaviour over geological time and remains essential for constraining long-term core dynamics (Suttie et al., 2025).

## 1.2 Spherical Harmonic Representation of the Geomagnetic Field

The measurements described in Section 1.1 provide the observational basis for studying the geomagnetic field. However, these measurements are unevenly distributed in space and time, and contain contributions from several different sources. In order to interpret them, it is therefore important to represent the field in a mathematical form that can be evaluated consistently at different locations and epochs. The starting point for this is Maxwell's equations, which describe the relationship between electric, $\boldsymbol{E}$ and magnetic, $\boldsymbol{B}$ fields, given by

$$\nabla \cdot \boldsymbol{E} = \frac{\rho}{\epsilon_0} \tag{1.1}$$

$$\nabla \times \boldsymbol{B} = \mu_0(\boldsymbol{J} + \epsilon_0 \frac{\partial \boldsymbol{E}}{\partial t}), \tag{1.2}$$

$$\nabla \cdot \boldsymbol{B} = 0, \tag{1.3}$$

$$\nabla \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t}. \tag{1.4}$$

where $t$ is time, $\mu_0$ is the permeability of free space, $\rho$ is the total electric charge density, $\boldsymbol{J}$ is the total electric current density, and $\epsilon_0$ is the permittivity of free space (Griffiths, 2017).

The geomagnetic field measured at and above Earth's surface contains contributions from several internal and external sources. At the frequencies relevant to the main geomagnetic field, the displacement current, $\epsilon_0 \partial \boldsymbol{E}/\partial t$, is negligible compared with the electric current density $\boldsymbol{J}$. Under this quasi-static approximation, Equation (1.2) reduces to

$$\nabla \times \boldsymbol{B} = \mu_0 \boldsymbol{J}. \tag{1.5}$$

In a source-free region, meaning a region containing no electric currents such that $\boldsymbol{J} = 0$, this becomes

$$\nabla \times \boldsymbol{B} = 0. \tag{1.6}$$

Since the magnetic field is curl-free in a source-free region, it can be expressed as the gradient of a scalar potential, $V$,

$$\boldsymbol{B} = -\nabla V. \tag{1.7}$$

Equation (1.3) states that there are no magnetic monopoles, meaning isolated magnetic poles

that act as sources or sinks of the magnetic field. Substituting Equation (1.7) into Equation (1.3) therefore gives Laplace's equation,

$$\nabla^2 V = 0. \tag{1.8}$$

The problem of describing the magnetic field outside the source region is therefore reduced to solving Laplace's equation for the scalar potential $V$, from which the magnetic field can be calculated using Equation (1.7). As the Earth is approximately spherical, it is convenient to solve Laplace's equation in spherical coordinates $(r, \theta, \phi)$, where $r$ is radius, $\theta$ is co-latitude, and $\phi$ is longitude. In spherical coordinates, Laplace's equation becomes

$$\nabla^2 V = \frac{1}{r^2}\frac{\partial}{\partial r}\left(r^2 \frac{\partial V}{\partial r}\right) + \frac{1}{r^2 \sin\theta}\frac{\partial}{\partial \theta}\left(\sin\theta \frac{\partial V}{\partial \theta}\right) + \frac{1}{r^2 \sin^2\theta}\frac{\partial^2 V}{\partial \phi^2} = 0. \tag{1.9}$$

Since $\boldsymbol{B} = -\nabla V$, the corresponding spherical components are

$$B_r = -\frac{\partial V}{\partial r}, \qquad B_\theta = -\frac{1}{r}\frac{\partial V}{\partial \theta}, \qquad B_\phi = -\frac{1}{r\sin\theta}\frac{\partial V}{\partial \phi}. \tag{1.10}$$

Following Merrill et al. (1996), the scalar potential may be separated into contributions from internal sources, located within the Earth, and external sources, associated with electric currents in the ionosphere and magnetosphere,

$$V(r, \theta, \phi, t) = V_{\text{int}}(r, \theta, \phi, t) + V_{\text{ext}}(r, \theta, \phi, t). \tag{1.11}$$

The field generated by internal sources can be written as

$$V_{\text{int}}(r, \theta, \phi, t) = a \sum_{n=1}^{N} \sum_{m=0}^{n} \left(\frac{a}{r}\right)^{n+1} \left[g_n^m(t)\cos(m\phi) + h_n^m(t)\sin(m\phi)\right] P_n^m(\cos\theta), \tag{1.12}$$

and the contribution from external sources as

$$V_{\text{ext}}(r, \theta, \phi, t) = a \sum_{n=1}^{N} \sum_{m=0}^{n} \left(\frac{r}{a}\right)^{n} \left[q_n^m(t)\cos(m\phi) + s_n^m(t)\sin(m\phi)\right] P_n^m(\cos\theta), \tag{1.13}$$

where $a$ is the radius of the Earth, $t$ is time, $n$ is the spherical harmonic degree, $m$ is the order, and $P_n^m(\cos\theta)$ are Schmidt semi-normalised associated Legendre functions. The coefficients $g_n^m(t)$, $h_n^m(t)$, $q_n^m(t)$, and $s_n^m(t)$ define the amplitude of each spherical harmonic contribution

to the field. The internal coefficients $g_n^m(t)$ and $h_n^m(t)$ are commonly referred to as Gauss coefficients, while $q_n^m(t)$ and $s_n^m(t)$ describe the external contribution. In this representation, each coefficient acts as a weight on a particular spatial pattern. The degree $n$ controls the characteristic angular length scale of the pattern, with higher degrees corresponding to finer, shorter-wavelength structure, while the order $m$ controls its longitudinal variation. The coefficients are estimated from magnetic-field observations by fitting the expansion up to a maximum degree $N$, which sets the smallest spatial scales represented in the model. The expansion begins at $n = 1$, since an $n = 0$ magnetic term would correspond to a monopole contribution, which is excluded by $\nabla \cdot \boldsymbol{B} = 0$. The explicit radial dependence of each harmonic degree also allows the geomagnetic field to be upward or downward continued to any radius.

The distribution of magnetic-field power across spatial scales can be described using the Mauersberger–Lowes spectrum (Mauersberger, 1956; Lowes, 1966). For the internal field, at a set time, this is calculated from the Gauss coefficients as

$$R_n(r,t) = (n+1)\left(\frac{a}{r}\right)^{2n+4} \sum_{m=0}^{n} \left[ \left(g_n^m(t)\right)^2 + \left(h_n^m(t)\right)^2 \right], \tag{1.14}$$

where $R_n(r,t)$ is the contribution to the mean square magnetic field from degree $n$, evaluated at radius $r$. Low degrees describe long-wavelength structure, while higher degrees describe progressively shorter-wavelength structure. The spectrum is therefore useful for comparing field models and for assessing how much structure is retained when the spherical harmonic expansion is truncated. It also provides a way to relate the mathematical model to the physical origin of the observed field.

## 1.3 Earth's Structure

The structure of the core–mantle boundary region is important for understanding the dynamics and evolution of the geomagnetic field. The present-day Earth comprises a solid inner core, a liquid outer core and a predominantly solid silicate mantle, with the chemically distinct crust forming the outermost layer. The inner core has a radius of 1225 km, while the outer core extends to the core–mantle boundary (CMB) at a radius of 3485 km (Dziewonski and Anderson, 1981). The core consists primarily of iron and nickel with a smaller proportion of light elements, while the mantle is composed mainly of silicate rock (see Hirose et al. (2021) for a full review of light elements in the core). The CMB marks a density increase of more than $4000\,\mathrm{kg\,m^{-3}}$ between the

lowermost mantle and outer core (Russell et al., 2024). The two regions also differ substantially in viscosity, and thermal and electrical conductivity, resulting in contrasting dynamics across the CMB. Understanding the exchange of heat and matter across this boundary is therefore important for linking the structure, dynamics and evolution of the core and mantle.

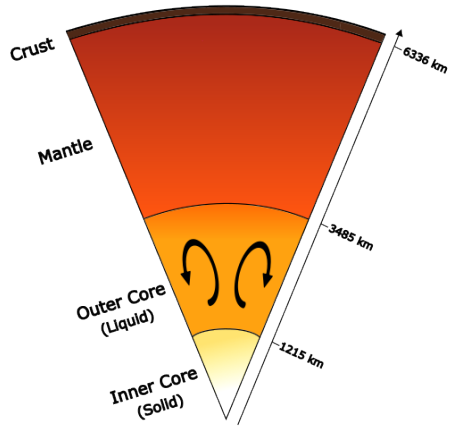



Figure 1.5: Schematic picture of the Earth, not to scale.

## 1.4 Sources of the Geomagnetic Field

At and above the Earth's surface, the measured geomagnetic field is the sum of several contributions from both internal and external sources, shown in the schematic in Figure 1.6. The largest contribution is the internal field, dominated by the core field generated through the interaction between the existing magnetic field and the motion of electrically conducting fluid in the outer core (Roberts, 2015). This is combined with the lithospheric field, produced by magnetised rocks in the crust and uppermost mantle. The remaining contributions arise from external current systems in the ionosphere and magnetosphere, driven by the interaction between charged particles and plasma from the Sun, and the Earth's magnetic environment. The movement of electrically conductive seawater also generates a magnetic field. These external fields are also important for studies of the internal field, as their time variation induces secondary currents in the solid Earth and oceans that obscure the core field signal. Although these sources differ in their physical origin, location, and typical spatial and temporal behaviour, they are not measured separately. A magnetic observation instead records their combined signal, so the different contributions must be separated in order to isolate the core field and study the processes responsible for secular variation (Olsen et al., 2010b; Hulot et al., 2015).

Separating the sources of the geomagnetic field is far from a trivial matter. The simplest

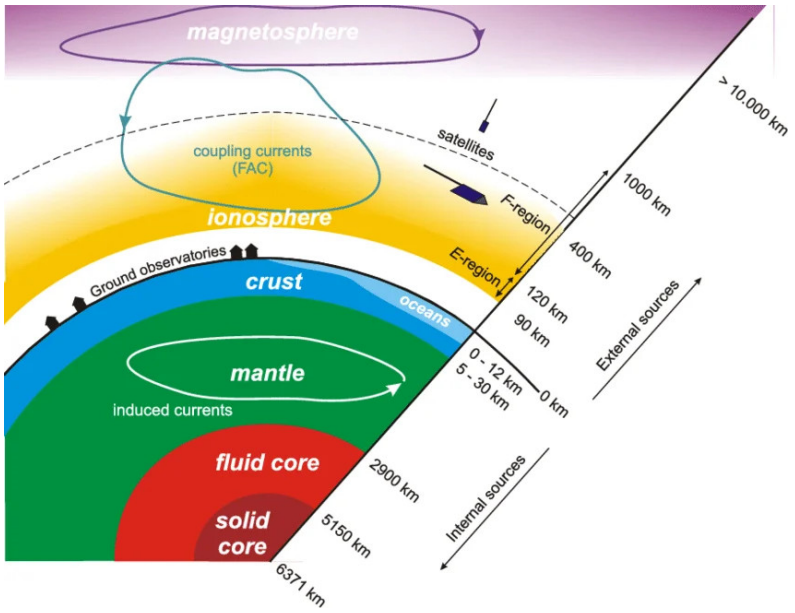



Figure 1.6: Sketch of the various sources contributing to the near-Earth magnetic field (Olsen et al., 2010b).

distinction is between fields of internal and external origin, following Gauss (1839). In the spherical harmonic representation described in Section 1.2, internal and external contributions have different radial dependences and are described by separate sets of coefficients. Fitting these contributions to observations of the magnetic-field components therefore allows the internal coefficients $(g_n^m, h_n^m)$ and external coefficients $(q_n^m, s_n^m)$ to be estimated separately (Olsen et al., 2010a). Alternative approaches aim to separate internal and external contributions using spatially localised basis functions (Chambodut et al., 2005; Stockmann et al., 2009; Beggan et al., 2013), which reduce sensitivity to uneven data coverage but do not fully resolve source separation, particularly for incomplete spatial sampling (Huang et al., 2026). Correlation-based methods separate geomagnetic field components by incorporating prior information about their spatial correlation structures within a Bayesian inversion, which also provides estimates of model uncertainty (Holschneider et al., 2016). However, source separation is not generally unique without additional constraints on the spatial distribution of the sources (Baratchart and Gerhards, 2017). Satellite gradient data provide further constraints because they are less sensitive to large-scale external-field fluctuations and improve the recovery of shorter-wavelength internal-field structure (Kotsiaros et al., 2015; Olsen et al., 2015). Gradient secular-variation data can also improve the spatial resolution of core-flow inversions relative to vector measurements (Whaler et al., 2022; Madsen et al., 2025).

Another way of separating the sources of the geomagnetic field is to consider the timescales on

which these sources vary, as shown in Figure 1.7. This approach requires time series observations of the geomagnetic field, which can be analysed using spectral methods to describe the power associated with variations at different frequencies, allowing the processes contributing to the field to be separated (Constable and Constable, 2004, 2023). The geomagnetic field exhibits variability over a wide range of timescales, from long-period internal processes such as geomagnetic reversals and secular variation of the core field, to high-frequency external variations generated by interactions between the solar wind, magnetosphere, and ionosphere. Spectral analyses show that the internal core field dominates at low frequencies, while external fields become dominant at periods of the order of the 11-year solar cycle and shorter (Constable and Constable, 2023). In practice, temporal separation exploits these differences either by averaging over time to reduce rapidly varying external fields (Constable, 2019), or by modelling time-dependent external contributions using indices such as Dst, which represent magnetospheric field variations (Maus and Weidelt, 2004).

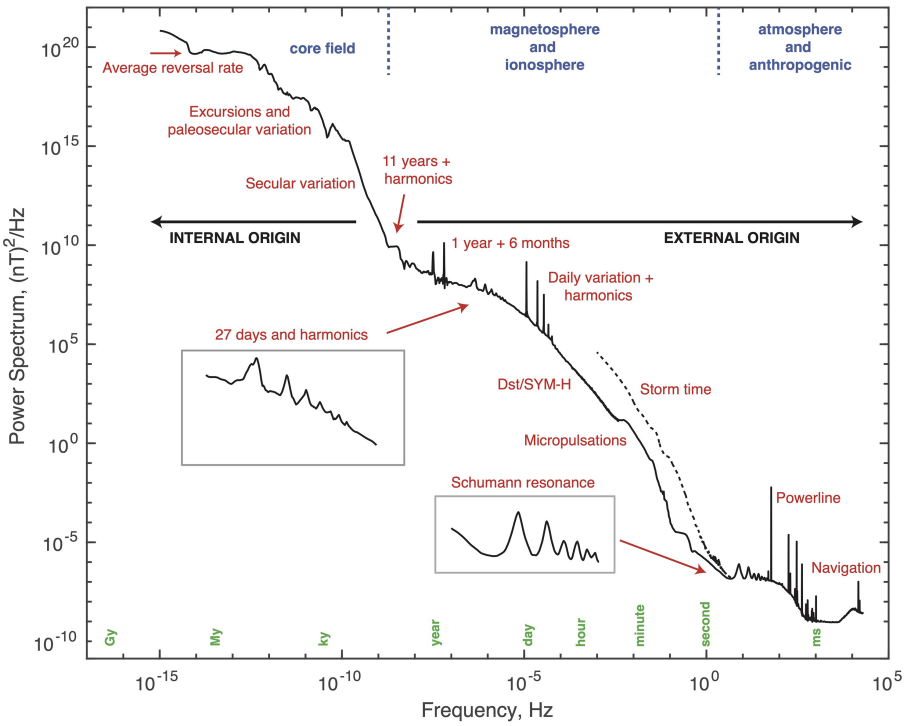



Figure 1.7: The Grand Spectrum of the geomagnetic field, showing dipole variations from global observations at long periods and horizontal field measurements from local data at shorter periods (Constable and Constable, 2023).

### 1.4.1 The Core Field

The core field is the dominant contribution to the geomagnetic field measured at the Earth's surface, and is generated by the internal, self-sustaining planetary dynamo, first identified as the largest source of the field by Bullard (1950). This dynamo is generated by the motion of electrically conducting liquid outer core, where heat and light elements released during the solidification of the inner core drive thermal and compositional convection (Deguen et al., 2013). The Earth's rotation also exerts Coriolis forces on this flow, helping to organise the convection and producing a field that is preferentially axial and dipolar at large scales (Cormier et al., 2022). As electrically conducting fluid moves through an existing magnetic field, it induces electric currents, which in turn generate a magnetic field. This continuous interaction between fluid motion, electric currents, and magnetic field converts kinetic energy into magnetic energy, maintaining the field against ohmic decay (Hollerbach, 1996). At the CMB, the magnetic field intensity is typically of the order of milliteslas, ranging from approximately 20,000–30,000 nT at low latitudes to 600,000–760,000 nT at high latitudes (Olsen et al., 2010b). Although the surface field is often described as a tilted dipole, this approximation only captures the largest-scale structure. Significant non-dipolar features are also present and these provide important information about the structure and evolution of the field generated in the core.

The radial core magnetic field for January 2026 is shown in Figure 1.8, with the field evaluated at the Earth's surface in the top panel and at the CMB in the bottom panel. Both maps are produced using CHAOS-8, a time-dependent geomagnetic field model constructed from satellite and ground-based observatory data, which provides estimates of the main field and secular variation (Kloss et al., 2026). At the Earth's surface, $B_r$ is dominated by the axial dipole, with mainly inward flux in the northern hemisphere and outward flux in the southern hemisphere (Glaßmeier et al., 2009; Hulot et al., 2015; Finlay et al., 2016). Regional departures from this simple structure are also visible, including the weak-field region associated with the South Atlantic Anomaly (Finlay et al., 2020; Meier et al., 2023; Gao et al., 2025). When the same field is downward continued to the CMB, the radial field reveals a more spatially complex structure than is visible at the Earth's surface (Gubbins and Bloxham, 1985). The large-scale dipole remains an important component of the field, but it is accompanied by intense flux patches and regions of weak or reversed flux at the CMB (Davies and Constable, 2017; Nilsson et al., 2020). In the present field, reversed-flux patches beneath the South Atlantic are

associated with the weak surface-field region known as the South Atlantic Anomaly, although the relationship between individual CMB patches and the surface intensity minimum is not one-to-one (Terra-Nova et al., 2017; Finlay et al., 2020; Gao et al., 2025).

The time dependence of the main field is described by its secular variation (SV), defined as the rate of change of the core field on timescales from months to centuries. In the spherical harmonic representation, the SV is given by the time derivatives of the Gauss coefficients, $\dot{g}_n^m(t)$ and $\dot{h}_n^m(t)$, and therefore describes how the spatial structure of the field evolves through time (Jackson and Finlay, 2015). Figure 1.9 shows the radial component of the SV in January 2026, evaluated at both the Earth's surface and the CMB using CHAOS-8 (Kloss et al., 2026). As with the main field, downward continuation to the CMB reveals shorter-wavelength structure that is attenuated at the Earth's surface.

The SV is not steady through time, as shown in the time–longitude plot in Figure 1.10. It includes long-term behaviour such as westward drift, axial dipole decay and the evolution of reversed-flux patches, together with more rapid changes known as geomagnetic jerks (Hulot et al., 2015; Jackson and Finlay, 2015). Jerks are abrupt changes in the trend of the SV and have been identified in both ground-based and satellite observations (Brown et al., 2013; Whaler et al., 2022; Madsen et al., 2025).

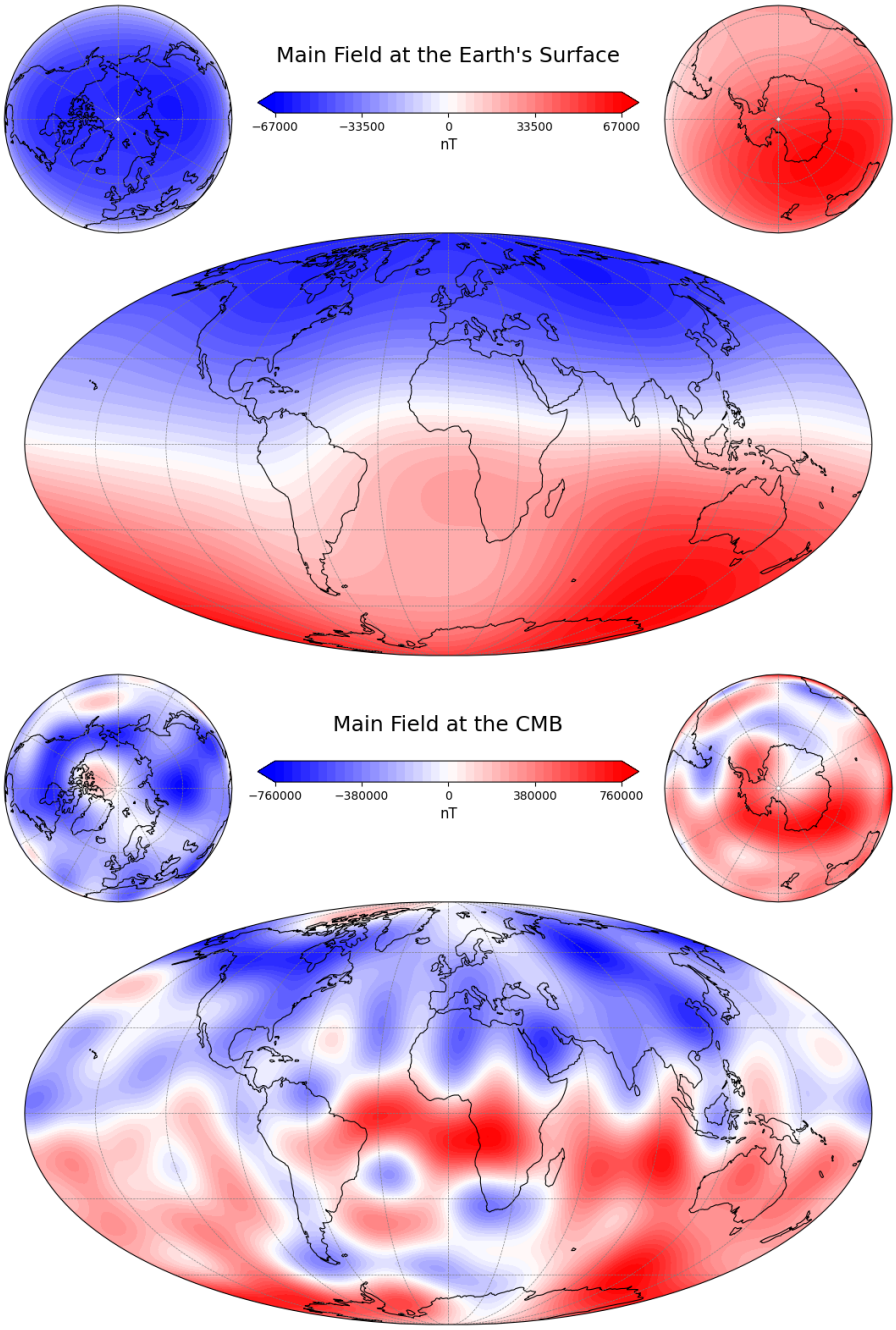



Figure 1.8: Radial component of the main field at the Earth's surface (top) and at the CMB (bottom), shown using the CHAOS-8 model (Kloss et al., 2026) (degree 12) for 1 January 2026.

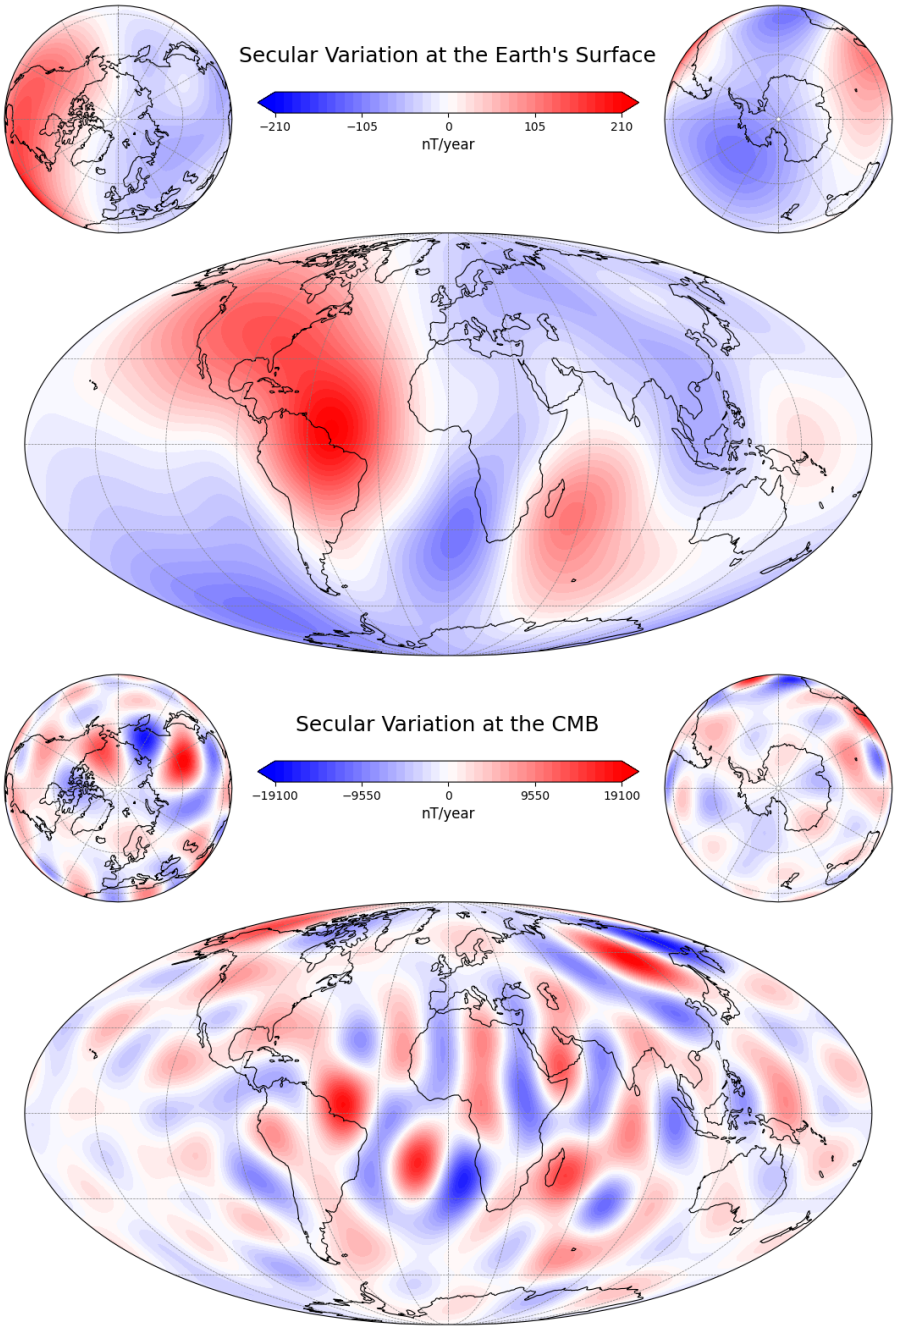



Figure 1.9: Radial component of the Secular Variation at the Earth's surface (top) and at the CMB (bottom), shown using the CHAOS-8 model (Kloss et al., 2026) (degree 12) for 1 January 2026.

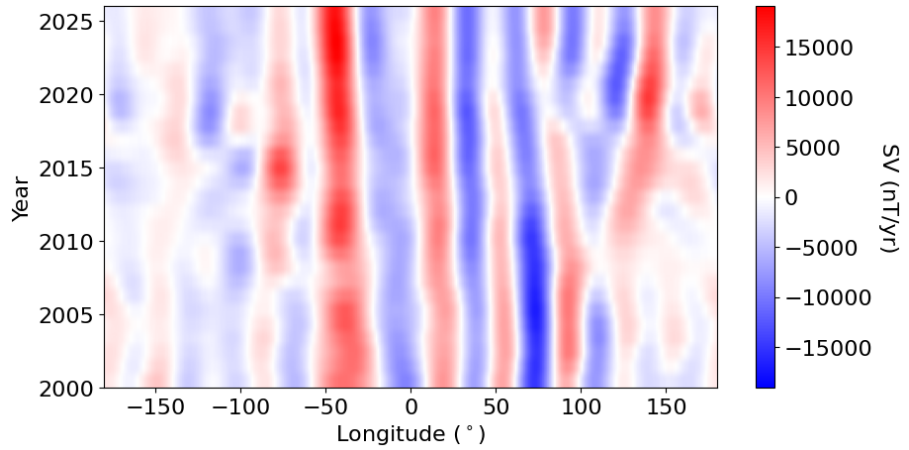



Figure 1.10: Time–longitude plot of the radial component of secular variation at the CMB along the equator, calculated from the CHAOS-8 model (Kloss et al., 2026) truncated at spherical harmonic degree 12, for 2000–2026.

### 1.4.2 The Lithospheric Field

The lithospheric, or crustal, field is produced by magnetised rocks in the crust and uppermost mantle, and represents the second major internal contribution to the geomagnetic field after the core field. It is much weaker than the main field at the Earth's surface, but contains more short-wavelength structure because it is produced by geological features that vary over smaller spatial scales, which themselves block the short-wavelength structure of the core field. The depth extent of the magnetic lithosphere is limited by the Curie temperature of its constituent magnetic minerals, above which permanent magnetisation is lost. This temperature depends on mineral composition, but is approximately 580 °C for magnetite. The thickness of the magnetic lithosphere therefore varies with mineralogy and the local thermal structure (Thébault et al., 2010).

Lithospheric magnetisation may be remanent or induced. Remanent magnetisation is acquired when magnetic minerals align with the ambient field during cooling, crystallisation or chemical alteration, and may be retained after the field in which it formed has changed. The pattern of geomagnetic polarity reversals is preserved in oceanic crust, where symmetric magnetic anomalies about mid-ocean ridges record alternating bands of normal and reversed polarity formed during seafloor spreading (Vine, 1966). These anomalies therefore provide a record of long-term variations in the geomagnetic field. In contrast, continental magnetic anomalies are generally less regular because they reflect greater variation in crustal age, composition, thickness and tectonic history.

Induced rock magnetisation is produced by the response of magnetic minerals to the present main field, and depends on the magnetic susceptibility of the material and the strength and direction of the inducing field (Thébault et al., 2010). Both remanent and induced magnetisation may contribute to the same lithospheric anomaly, with their relative importance depending on the magnetic properties and geological history of the source region. Induced rock magnetisation is distinct from the magnetic fields produced by electrical currents induced in the conducting Earth and oceans by time-varying external fields, which are discussed in the following section.

Lithospheric field models are used to map magnetic anomalies and relate them to crustal and upper-mantle structure. Global compilations such as the World Digital Magnetic Anomaly Map (WDMAM), the first version of which is shown in Figure 1.11, combine satellite, ma-

rine, aeromagnetic and ground-based measurements to provide a global view of lithospheric magnetic variations (Thébault et al., 2010). The most recent version is found in Lesur et al. (2016). Satellite observations constrain the longer-wavelength part of the lithospheric field, while shorter-wavelength structure is better resolved by near-surface marine and aeromagnetic surveys (Olsen et al., 2017). The coverage and resolution of these measurements remain spatially uneven, particularly across parts of the oceans and continents (Lesur et al., 2016). Separating the lithospheric field from the small-scale core field is a significant challenge because their contributions overlap at intermediate spherical harmonic degrees. This makes it difficult to identify the 'knee' in the Lowes spectrum, shown in Figure 1.12, which marks the transition from the core field at larger spatial scales to the lithospheric field at smaller spatial scales. Consequently, incomplete separation of these fields limits the maximum spherical harmonic degree to which the core field can be reliably resolved (Otzen et al., 2024).

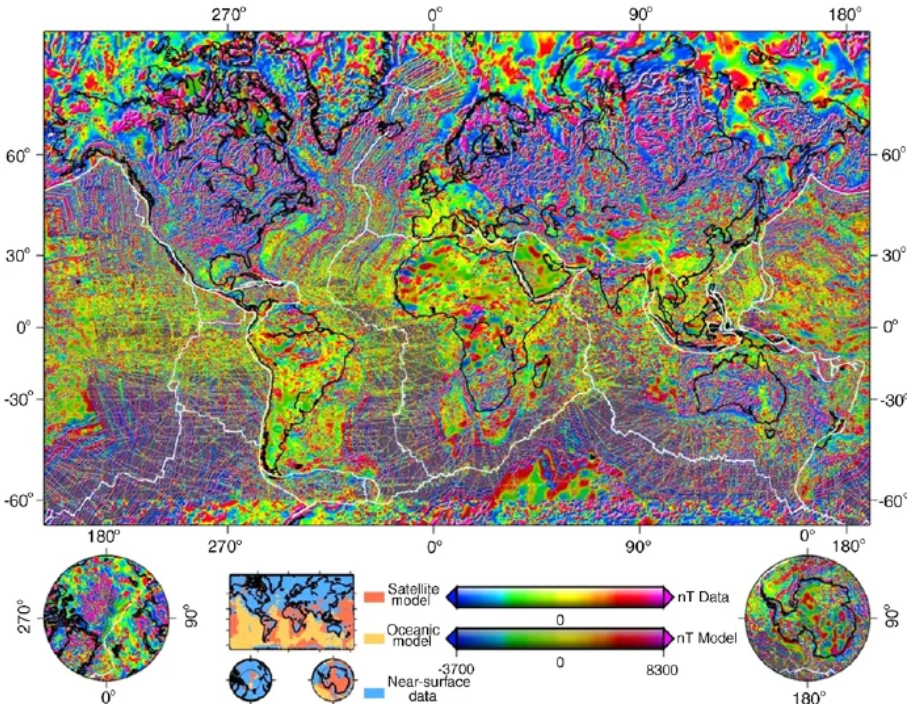



Figure 1.11: The first version of the World Digital Magnetic Anomaly Map (WDMAM) (Thébault et al., 2010), showing magnetic anomalies, which are local deviations from the expected geomagnetic field caused by variations in the magnetisation of crustal rocks. The thumbnail at the bottom illustrates the geographical distribution of the various information contained in the map.

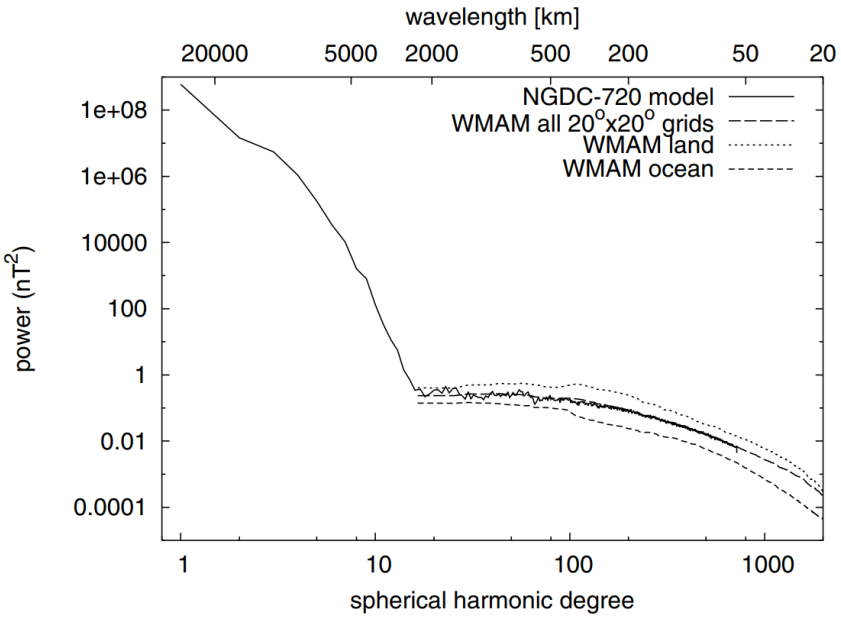



Figure 1.12: Geomagnetic power spectra for the NGDC-720 model and subgrids of the NGDC World Magnetic Anomaly Map. The 'knee' in the spectrum marks the transition between the predominantly core-generated field at lower spherical harmonic degrees and the predominantly lithospheric field at higher degrees. Taken from Maus (2008)

### 1.4.3 The External Field

The external field is produced by electric currents outside the solid Earth, mainly in the ionosphere and magnetosphere, a scehamtic for which is shown in Figure 1.13. These currents are driven by the interaction between the solar wind and the Earth's magnetic environment, and vary much more rapidly than the core and lithospheric fields (Baumjohann and Nakamura, 2007; Olsen et al., 2010b). During geomagnetically quiet periods, the external field is usually much weaker than the main field, with ionospheric and magnetospheric contributions of the order of tens of nanotesla at the Earth's surface (Hulot et al., 2015; Olsen and Stolle, 2017). During disturbed periods, such as geomagnetic storms and substorms, these current systems intensify and produce larger magnetic variations over timescales from seconds to days (Kauristie et al., 2017).

Time-varying external fields also induce electrical currents in the conducting Earth and oceans. These currents produce secondary magnetic fields that are internal in origin, but are driven by external-field variations (Maus and Weidelt, 2004; Olsen et al., 2010b). They are therefore distinct from the induced magnetisation of lithospheric rocks described above. Comprehensive geomagnetic field models commonly represent the primary external field together with its internally induced counterpart. For example, CHAOS-8 includes contributions from magnetospheric currents and the associated induced response when separating these signals from the core and lithospheric fields (Kloss et al., 2026).

External and induced fields must therefore be accounted for when modelling the internal field. This is particularly important for satellite observations at high latitudes, where polar ionospheric currents, auroral electrojets and field-aligned currents produce magnetic signals that can contaminate estimates of the core and lithospheric fields (Finlay et al., 2017; Kloss et al., 2023). Internal-field models commonly use quiet-time data selection and explicit parameterisations of external and induced contributions to reduce this contamination (Kauristie et al., 2017; Finlay et al., 2017; Kloss et al., 2026). However, the remaining contamination is spatially structured and may be concentrated in particular regions, especially around the poles.

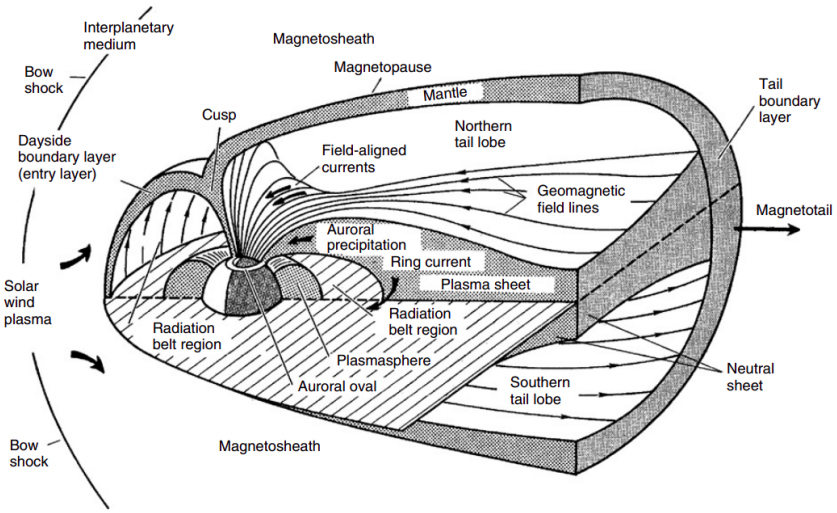



Figure 1.13: Schematic of the Earth's magnetosphere and the principal external current systems (Kelley and Heelis, 1989).

## 1.5 Modelling and Predicting the Geomagnetic Field

Geomagnetic field models combine measurements from satellites, ground-based observatories and surveys to provide a consistent representation of the field across locations and epochs. These models are required because observations are unevenly distributed in space and time and contain overlapping contributions from several sources. They provide reference fields for navigation and geomagnetic surveying, while also allowing the spatial structure and temporal evolution of the internal field to be studied. Since the core field changes through time, models may also be used to estimate its secular variation and forecast its evolution beyond the period covered by observations.

Most geomagnetic field models represent the magnetic potential using the spherical harmonic expansion introduced in Section 1.2, but differ in how the coefficients and their time dependence are estimated. One common approach is regularised inversion, which minimises the misfit to the observations alongside a constraint on the spatial or temporal complexity of the solution. The misfit may be defined using an $L_2$ norm or a robust, iteratively re-weighted $L_1$ norm that reduces the influence of outliers (Beggan and Whaler, 2008). Regularisation may constrain small-scale structure or changes between successive epochs (Beggan and Whaler, 2008; Kloss and Finlay, 2019; Whaler et al., 2016, 2022). This stabilises poorly constrained inversions, but the recovered structure depends on the chosen norm, constraint and damping strength. Weaker regularisation permits a closer fit to the data, but may introduce poorly resolved structure, whereas stronger regularisation produces smoother solutions but may suppress genuine variations. Other approaches include data assimilation, which combines geomagnetic observations with prior information about the spatial, temporal or dynamical evolution of the field (Fournier et al., 2010; Baerenzung et al., 2020).

The following sections focus on the International Geomagnetic Reference Field and the CHAOS family of models, as these are the most relevant to this thesis, although several other geomagnetic field models exist. These include COV-OBS, a time-dependent field model that uses stochastic prior information derived from the temporal spectra of geomagnetic observations and provides estimates of both the field and its associated uncertainties (Gillet et al., 2015). Kalmag similarly uses a Bayesian framework, but constructs the model sequentially by combining a Kalman filter with a smoothing algorithm, allowing the complete dataset to inform the field estimate at

each epoch. It jointly estimates the core, lithospheric, magnetospheric and induced fields, and provides mean solutions together with their associated uncertainties (Baerenzung et al., 2020).

### 1.5.1 The International Geomagnetic Reference Field

The International Geomagnetic Reference Field (IGRF) is a mathematical representation of the large-scale, time-varying internal magnetic field. It is produced by the *Geomagnetic Field Modelling Working Group* of the *International Association of Geomagnetism and Aeronomy* and is revised every five years. The model consists of spherical harmonic coefficients that can be used to calculate the three vector components of the magnetic field above the CMB. It provides a standard reference field for applications including navigation, geomagnetic surveying and studies of the evolution of the core field (Alken et al., 2021b).

Each generation of the IGRF is constructed from candidate models submitted by international research groups. These candidates are evaluated and combined to produce agreed main-field and predictive secular-variation models. IGRF-13 provides coefficients spanning 1900–2025 and contains a definitive main-field model for epoch 2015.0, a main-field model for epoch 2020.0 and a predictive linear secular-variation model for 2020.0–2025.0 (Alken et al., 2021b,a). Eleven candidate models were submitted for epoch 2015.0, twelve for epoch 2020.0 and fourteen for the 2020–2025 secular-variation forecast. The final main-field models were calculated from the median of the candidate Gauss coefficients, while the predictive secular-variation model was obtained using a robust Huber-weighted combination (Alken et al., 2021a). The latest iteration, IGRF-14, extends the model series to 2030, providing an updated main-field model for epoch 2025.0 and a predictive secular-variation model for 2025.0–2030.0. Its coefficients were finalised in November 2024 following the evaluation and combination of candidate models submitted by international research groups (Beggan et al., 2026).

The IGRF candidate process provides a framework for comparing a wide range of approaches to modelling and forecasting the geomagnetic field. Candidate methodologies for secular variation have included constructing models from ground and satellite measurements (Alken et al., 2021a; Finlay et al., 2020; Rother et al., 2021; Huder et al., 2020), data assimilation techniques (Fournier et al., 2021; Minami et al., 2020; Sanchez et al., 2020), Kalman-filter based models (Tangborn et al., 2021; Baerenzung et al., 2020), interpolation (Petrov and Bondar, 2021), diffusion based models (Metman et al., 2020), bootstrapping (Pavón-Carrasco et al., 2020) and spectral analysis

(Wardinski et al., 2020).

### 1.5.2 The CHAOS Field Models

The CHAOS series provides time-dependent spherical harmonic based models of the near-Earth geomagnetic field constructed from satellite and ground-based observatory measurements, describing the continuous evolution of the field (Olsen et al., 2006; Finlay et al., 2020). The temporal evolution of the internal field is represented using B-splines, while regularisation limits poorly constrained changes between observations. CHAOS models are updated as new measurements become available and have provided parent models for candidate submissions to the IGRF.

#### CHAOS-7

CHAOS-7 describes the near-Earth geomagnetic field between 1999 and 2024 using measurements from the Swarm, CryoSat-2, CHAMP, SAC-C and Ørsted satellites, together with ground-based observatory data (Finlay et al., 2020). The model contains a time-dependent internal field to spherical harmonic degree 20, a static internal field above degree 20, and representations of the magnetospheric field and its internally induced counterpart. Its parameters are estimated using an iteratively reweighted regularised least-squares procedure applied to data selected under strict geomagnetically quiet-time criteria. The internal-field coefficients are represented using B-splines and constrained by temporal regularisation, providing continuous estimates of the main field and secular variation for use in core-flow inversion.

#### CHAOS-8

CHAOS-8 extends the model series to the present day while retaining the principal structure of CHAOS-7 (Kloss et al., 2026). It contains a time-dependent internal field to spherical harmonic degree 20, a static internal field at higher degrees, and representations of the magnetospheric field and its induced counterpart. An initial version of CHAOS-8 provided the parent model for the DTU candidates submitted for IGRF-14.

The main differences between the modelling approaches of CHAOS-7 and CHAOS-8 are the treatment of ionospheric contamination and the temporal regularisation of the internal field. CHAOS-8 includes an additional model of ionospheric contributions that may remain significant under dark and geomagnetically quiet conditions (Kloss et al., 2026). It uses temporal

covariances extracted from geodynamo simulations (Aubert, 2023; Gillet et al., 2024) (see Section 1.10) to regularise the time-dependent internal field throughout the model interval, while retaining a penalty on the second time derivative of the radial field at the model endpoints. In contrast, CHAOS-7 penalised the third time derivative of the radial field throughout the model interval, in addition to the endpoint regularisation. This change allows more rapid variations to be recovered at shorter spatial scales, producing stronger and more clearly defined sub-decadal secular-variation features in CHAOS-8 than in CHAOS-7, with time–longitude plots showing the differences in Figure 1.14. These include stationary and propagating patterns that have been interpreted as possible signatures of hydromagnetic waves (Gillet et al., 2022; Whaler et al., 2022; Grüne et al., 2025), although their recovery remains limited by the spatial resolution of the core field.

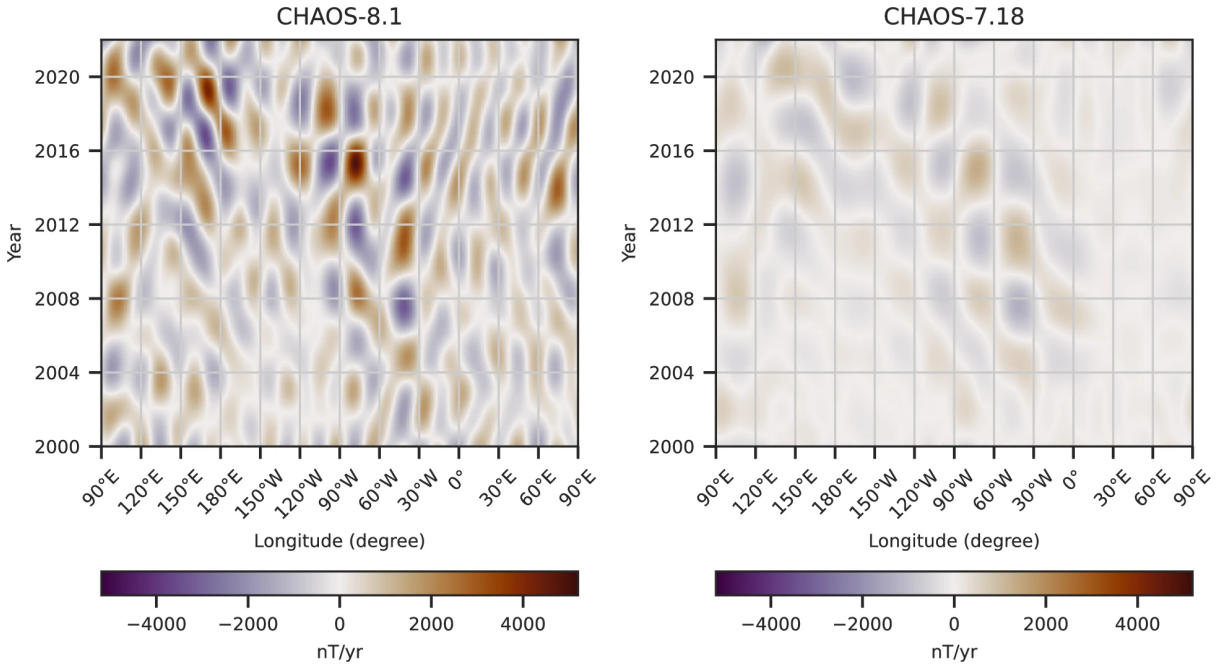



Figure 1.14: Time–longitude plots of the radial secular variation along the equator at the CMB between 2000 and 2022, high-pass filtered using a 9-year cut-off period. Results are shown for CHAOS-8.1 (left) and CHAOS-7.18 (right). Taken from Kloss et al. (2026)

## 1.6 Relating Secular Variation to Core Flow

As discussed in Section 1.4.1, the magnetic field within the Earth is generated within the Earth's outer core. This evolution is governed by the induction equation, which describes the advection of the magnetic field by the core flow and its diffusion through the fluid. Ohm's law for a conductor moving with velocity $\boldsymbol{u}$ relates the total electric current density, $\boldsymbol{J}$, to the electric field, $\boldsymbol{E}$, and is given by

$$\boldsymbol{J} = \sigma(\boldsymbol{E} + \boldsymbol{u} \times \boldsymbol{B}), \tag{1.15}$$

where $\sigma$ is the conductivity and $\boldsymbol{B}$ is the magnetic field. Substituting Ohm's law into Equation (1.2) and neglecting the displacement current, such that $\nabla \times \boldsymbol{B} = \mu_0 \boldsymbol{J}$, gives

$$\nabla \times \boldsymbol{B} = \mu_0 \sigma(\boldsymbol{E} + \boldsymbol{u} \times \boldsymbol{B}). \tag{1.16}$$

Taking the curl of Equation (1.16) and rearranging gives

$$\nabla(\nabla \cdot \boldsymbol{B}) - \nabla^2 \boldsymbol{B} = \mu_0 \sigma(\nabla \times \boldsymbol{E} + \nabla \times (\boldsymbol{u} \times \boldsymbol{B})), \tag{1.17}$$

Substituting Equations (1.3) and (1.4) into this expression and rearranging then gives the induction equation

$$\frac{\partial \boldsymbol{B}}{\partial t} = \nabla \times (\boldsymbol{u} \times \boldsymbol{B}) + \eta \nabla^2 \boldsymbol{B}, \tag{1.18}$$

where $\eta = \frac{1}{\mu_0 \sigma}$. The change in the magnetic field, known as secular variation (SV), therefore results from the combined effects of induction by the motion of the electrically conducting fluid, called advection, $\nabla \times (\boldsymbol{u} \times \boldsymbol{B})$, and ohmic diffusion, $\eta \nabla^2 \boldsymbol{B}$ (Roberts, 2015). Assuming that both terms have a characteristic length scale $L$, such that $\nabla \sim 1/L$, and introducing a characteristic velocity $U$, such that $\boldsymbol{u} \sim U$, the advection term has characteristic magnitude

$$\nabla \times (\boldsymbol{u} \times \boldsymbol{B}) \sim \frac{UB}{L}, \tag{1.19}$$

while the diffusion term has characteristic magnitude

$$\eta \nabla^2 \boldsymbol{B} \sim \frac{\eta B}{L^2}. \tag{1.20}$$

The ratio of the advection term to the diffusion term defines the magnetic Reynolds number, $R_m$,

$$R_m = \frac{UL}{\eta}. \tag{1.21}$$

Using typical values for the core (Holme, 2015), $U \cong 5 \times 10^{-4}\,\mathrm{m\,s^{-1}}$, $L \cong 1000\,\mathrm{km}$ and $\eta = 1\,\mathrm{m^2\,s^{-1}}$, gives $R_m \approx 500$. This indicates that advection dominates diffusion at this length scale, allowing the diffusion term in Equation (1.18) to be neglected. The induction equation then reduces to

$$\frac{\partial \boldsymbol{B}}{\partial t} = \nabla \times (\boldsymbol{u} \times \boldsymbol{B}), \tag{1.22}$$

which is known as the *frozen-flux approximation* (Roberts and Scott, 1965). Under this approximation, the magnetic field is advected by the motion of the electrically conducting fluid in the outer core and acts as a tracer of that motion (Roberts and Scott, 1965; Bloxham and Jackson, 1991; Holme, 2015; Jackson and Finlay, 2015). Changes in the magnetic field can therefore be related to fluid motion near the top of the core.

Approximating the CMB as a spherical material boundary implies that no fluid crosses it, such that $u_r = 0$ (Roberts and Scott, 1965; Holme, 2015). Under these assumptions, the radial component of the induction equation becomes

$$\frac{\partial B_r}{\partial t} + (\boldsymbol{u}_h \cdot \nabla_h) B_r + B_r \nabla_h \cdot \boldsymbol{u}_h = 0, \tag{1.23}$$

where $B_r$ is the radial magnetic field, $\boldsymbol{u}_h$ is the horizontal flow, and $\nabla_h$ denotes the horizontal gradient at the CMB. Since $B_r$ can be inferred from observations at the Earth's surface, this equation provides an observational link between secular variation and fluid motion at the top of the core (Holme, 2015; Jackson and Finlay, 2015).

The validity of the frozen-flux approximation remains uncertain. Although $R_m \approx 500$ suggests that diffusion is small at large scales, shorter radial length scales may exist near the CMB due to hydromagnetic boundary layers, increasing its relative importance (Bloxham and Jackson, 1991; Holme, 2015). Several studies have therefore considered how diffusion contributes to secular variation and how it may be incorporated into core-flow models (Gubbins, 1996; Amit and Christensen, 2008). Purely diffusive models can reproduce several large-scale features of decadal field evolution, including westward drift, movement of the North Magnetic Pole and the

emergence of reversed-flux patches, but do not capture rapid changes such as geomagnetic jerks (Metman et al., 2019). Diffusion alone also cannot sustain the magnetic field over geological timescales. Neither frozen flux nor pure diffusion therefore provides a complete description of secular variation, which is likely to result from the combined effects of fluid motion and magnetic diffusion (Metman et al., 2019).

Despite these limitations, the radial frozen-flux equation remains the basis of most core-flow studies. Given estimates of $B_r$ and its secular variation, $\partial B_r/\partial t$, it can be used to infer the large-scale horizontal motion responsible for the observed secular variation (Roberts and Scott, 1965; Bloxham and Jackson, 1991; Holme, 2015; Whaler, 1986; Jackson and Finlay, 2015).

## 1.7 Poloidal–Toroidal Decomposition

Under the Boussinesq approximation, density variations are neglected except where they contribute to buoyancy, so the outer-core fluid is treated as incompressible. Conservation of mass therefore requires the flow to be divergence-free, such that

$$\nabla \cdot \boldsymbol{u} = 0. \tag{1.24}$$

Together with Equation (1.3), which expresses the absence of magnetic monopoles, this means that both the flow and magnetic field are divergence-free (Holme, 2015). Since the outer-core flow is divergence-free, it may be represented in terms of poloidal and toroidal components. In spherical geometry, the velocity field can be written as

$$\boldsymbol{u} = \nabla \times (t\boldsymbol{r}) + \nabla \times \nabla \times (p\boldsymbol{r}), \tag{1.25}$$

where $t$ and $p$ are the toroidal and poloidal scalar functions, respectively, and $\boldsymbol{r}$ is the position vector (Holme, 2015). The toroidal component is given by

$$\boldsymbol{u}_T = \nabla \times (t\boldsymbol{r}), \tag{1.26}$$

and has no radial component. The poloidal component is given by

$$\boldsymbol{u}_P = \nabla \times \nabla \times (p\boldsymbol{r}), \tag{1.27}$$

and contains both radial and horizontal components. At the CMB, the impenetrable boundary condition requires the radial velocity to vanish. The decomposition therefore allows both toroidal and poloidal horizontal motion at the top of the core and forms the basis of many core-flow inversion methods.

## 1.8 Flow Non-Uniqueness

The radial frozen-flux equation provides only a single equation for the two horizontal components of flow. Additional assumptions are therefore required to reduce this non-uniqueness and obtain meaningful flow estimates (Backus, 1968; Bloxham and Jackson, 1991; Holme, 2015). These assumptions are commonly motivated by the force balance governing fluid motion within the outer core. Regularisation is used separately to favour a large-scale solution by penalising the magnitude or small-scale structure of the flow. Although this produces a numerically unique solution, it does not resolve the underlying physical non-uniqueness because other, slightly more complex flows may fit the observations within their uncertainties. The regularised solution is therefore one preferred member of a wider set of admissible flows and may not represent the true flow if the imposed regularisation is too restrictive (Holme, 2015). Additionally, regularisation and the assumption that the flow is dominated by large scales may blur separate regional features into a larger global structure. A key equation describing the dynamics of the outer core is the Navier–Stokes equation in a rotating reference frame. Under the Boussinesq approximation, the momentum equation becomes

$$\rho_0 \left( \frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u} \cdot \nabla \boldsymbol{u} + 2\boldsymbol{\Omega} \times \boldsymbol{u} \right) = -\nabla p + \rho' \boldsymbol{g} + \boldsymbol{J} \times \boldsymbol{B} + \rho_0 \nu \nabla^2 \boldsymbol{u}, \tag{1.28}$$

where $\rho_0$ is the hydrostatic density, $\rho'$ is the density peturbation, $\boldsymbol{\Omega}$ is the Earth's rotation vector, $p$ is the non-hydrostatic pressure, $\boldsymbol{g}$ is the acceleration due to gravity, $\boldsymbol{J}$ is the current density, and $\nu$ is the kinematic viscosity (Holme, 2015). The terms in Equation (1.28) represent inertia, the Coriolis force, the pressure gradient, buoyancy, the Lorentz force and viscosity, respectively.

Figure 1.15 shows two core-flow models derived from the same magnetic satellite data and inversion methodology, but using different flow assumptions (Holme and Olsen, 2006). Both recover similar large-scale features, including a gyre connecting westward flow beneath the Atlantic with circulation in the polar regions. However, the smaller-scale structure differs

between the models, particularly beneath the Pacific, demonstrating the sensitivity of inferred core flow to the imposed assumptions.

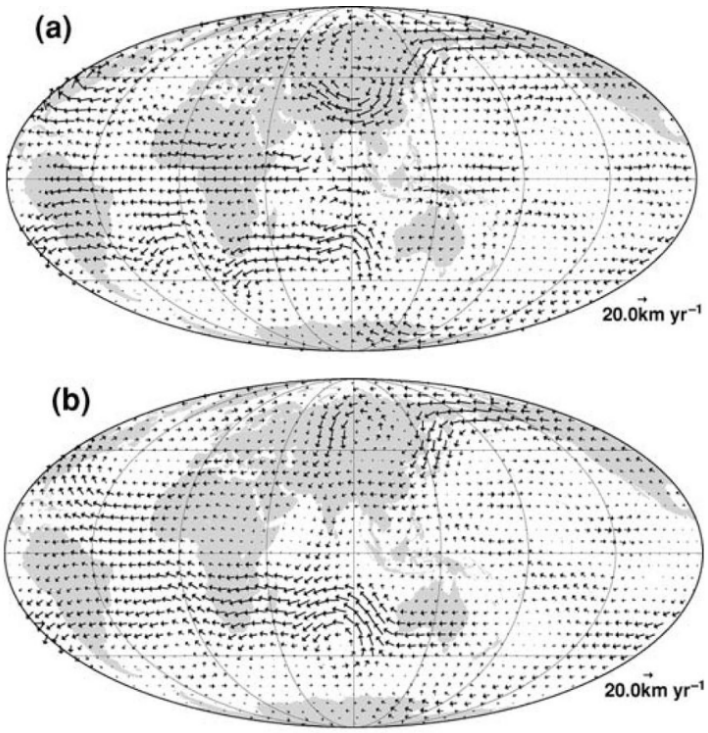



Figure 1.15: Tangentially geostrophic flow (a) and toroidal flow (b) derived from magnetic satellite data for 1999–2004. Adapted from Holme and Olsen (2006).

### 1.8.1 Tangentially Geostrophic Flow

Tangential geostrophy is a commonly used constraint in core-flow modelling and is based on the force balance near the top of the core (Le Mouël, 1984; Bloxham and Jackson, 1991). It assumes that the dominant forces are the Coriolis force, pressure gradient and buoyancy, while the Lorentz, inertial and viscous forces are neglected. Under this approximation, Equation (1.28) reduces to

$$2\rho_0 \left(\boldsymbol{\Omega} \times \boldsymbol{u}\right) = -\nabla p + \rho' \boldsymbol{g}. \tag{1.29}$$

Taking the curl removes the pressure term and gives the thermal wind equation

$$2\rho_0 \left(\boldsymbol{\Omega} \cdot \nabla\right) \boldsymbol{u} = \boldsymbol{g} \times \nabla \rho'. \tag{1.30}$$

Assuming that gravity is radial, the radial component of Equation (1.30) gives the tangentially geostrophic constraint

$$\nabla_h \cdot \left(\boldsymbol{u}_h \cos\theta\right) = 0, \tag{1.31}$$

where $\boldsymbol{u}_h$ is the horizontal flow and $\theta$ is co-latitude (Le Mouël, 1984; Amit and Pais, 2013).

Using

$$\boldsymbol{u}_h B_r = \left(\boldsymbol{u}_h \cos\theta\right) \left(\frac{B_r}{\cos\theta}\right), \tag{1.32}$$

the horizontal divergence term in the radial frozen-flux equation (1.23) can be expanded as

$$\nabla_h \cdot \left(\boldsymbol{u}_h B_r\right) = \frac{B_r}{\cos\theta} \nabla_h \cdot \left(\boldsymbol{u}_h \cos\theta\right) + \cos\theta\, \boldsymbol{u}_h \cdot \nabla_h \left(\frac{B_r}{\cos\theta}\right). \tag{1.33}$$

The first term on the right-hand side vanishes under the tangentially geostrophic constraint in Equation (1.31). The radial frozen-flux equation therefore becomes

$$\frac{\partial B_r}{\partial t} + \cos\theta\, \boldsymbol{u}_h \cdot \nabla_h \left(\frac{B_r}{\cos\theta}\right) = 0. \tag{1.34}$$

This reduces the non-uniqueness of the inferred flow but does not remove it. Flow remains undetermined along closed contours of $B_r/\cos\theta$, which bound regions known as ambiguous patches and typically cover approximately 40% of the CMB (Backus and Le Mouël, 1986; Bloxham and Jackson, 1991).

### 1.8.2 Other Flow Constraints

Several additional constraints have been proposed to reduce the non-uniqueness of core-flow inversion. One of these is the toroidal-flow assumption, in which radial motion is suppressed throughout the free stream of the outer core, leaving only the toroidal component of the velocity field (Holme, 2015). This reduces, but does not remove, the non-uniqueness of the inversion, since flow parallel to contours of constant $B_r$ produces no secular variation and therefore remains unconstrained (Backus, 1968). Observational support for the assumption was provided by Whaler (1980), who found that local extrema of $B_r$ were broadly consistent with non-divergent horizontal flow. However, uncertainties introduced by downward continuation prevent a conclusive test at the CMB. Accounting for these uncertainties, Whaler and Holme (2007) showed that standard core-flow inversions remain consistent with the horizontal convergence and horizontal divergence implied by the frozen-flux equation. Although toroidal-flow models have been successfully constructed (Whaler, 1986; Bloxham and Jackson, 1992), they generally provide a poorer fit to the observed secular variation than models containing both toroidal and poloidal components (Holme, 2015).

Quasi-geostrophic (QG) flow assumes that the horizontal pressure gradient and Coriolis forces remain close to balance throughout much of the outer core (Hollerbach and Gubbins, 2007; Gillet et al., 2012). Amit and Olson (2004) also introduced columnar and helical flow constraints. The former assumes that convection is organised into columns aligned with the Earth's rotation axis, while the latter imposes a relationship between horizontal flow divergence and vorticity motivated by numerical dynamo simulations. Although these assumptions further reduce the non-uniqueness of the inversion, they are generally less widely used than the tangentially geostrophic approximation (Amit and Olson, 2004; Holme, 2015).

## 1.9 Comparison of Regional and Global Flow

Core-surface flow is commonly inferred using global spherical harmonic representations of the magnetic field, secular variation and velocity field. These methods have identified several persistent large-scale features, most notably an eccentric, anticyclonic gyre beneath the Atlantic hemisphere (Pais and Jault, 2008; Kloss and Finlay, 2019; Gillet et al., 2022). More localised features have also been recovered, including the accelerating high-latitude jet beneath the Bering Strait (Livermore et al., 2017; Finlay et al., 2023), changes in equatorial azimuthal flow beneath Indonesia and Central America (Whaler et al., 2022; Li et al., 2024), and spatially localised changes associated with geomagnetic jerks (Brown et al., 2013). However, because spherical harmonics are a global basis, the flow at any point depends on information from across the CMB.

Regional inversions provide an independent means of testing whether features identified in global models are supported by the observations, or are influenced by the global representation and imposed large-scale regularisation. For example, a global inversion may connect westward flow beneath the Atlantic with motion towards the polar regions to produce a planetary-scale gyre. A regional inversion can test whether the individual features of this circulation are recovered when only information from that region is used. Agreement between the global and regional approaches would provide evidence that the gyre reflects the underlying magnetic-field data, whereas disagreement could indicate sensitivity to the global representation. Regional methods can similarly determine whether smaller-scale jets, vortices, and changes in flow direction are robust.

The difference between local and regional inversions is important when considering the influence of local unmodelled signals, such as external fields in the auroral zone. In a regional inversion, the same point depends only on information within the chosen region, preventing areas with greater uncertainty or unmodelled signals from affecting the local solution. This is particularly important because satellite measurements are more susceptible to contamination from external fields in auroral regions, while ground-based observatories are distributed unevenly. Regional inversions can therefore test the reliability of features in poorly constrained areas, while ensuring that flow estimates elsewhere are not influenced by uncertain polar data. This approach may become increasingly useful with enhanced regional sampling from the low-inclination MSS-1

mission and additional magnetic observations from proposed satellite constellations such as NanoMagSat (Jiang et al., 2024; Deconinck et al., 2025). Localised estimates from geomagnetic virtual observatories or Subtractive Optimally Localised Averages could further reduce reliance on an intermediate global field model (Mandea and Olsen, 2006; Hammer et al., 2021; Hammer and Finlay, 2018).

Previous studies have demonstrated the difficulties associated with recovering core flow on a regional scale. Rogers (2022) investigated regional variations using spherical Slepian functions, which improved the separation of secular variation at the Earth's surface compared with spherical harmonics. However, downward continuation amplified short-wavelength structure, while leakage across the regional boundaries and sensitivity to the regularisation limited the separation of secular variation and flow at the CMB. A pointwise inversion developed by Schwaiger et al. (2023) recovered the principal features found in global flow models, but required additional smoothing to suppress poorly resolved small-scale structure and was sensitive to the prior used to reduce the non-uniqueness. These studies highlight the challenges of developing a regional method that preserves local structure without introducing boundary leakage or poorly constrained small-scale flow.

## 1.10 Numerical Simulations

Numerical simulations provide an important means of investigating the dynamics of Earth's core and the geodynamo, complementing the measurement-driven inversion approaches in previous sections. They usually solve for the dimensionless fluid flow $\boldsymbol{u}^*$, magnetic field $\boldsymbol{B}^*$ and density anomalies $C^*$ throughout the core at spatial and temporal resolutions beyond those available from observations, by solving for the underlying physics. This provides access to processes that cannot be observed directly, including the small-scale core field obscured by the lithospheric field. Simulations also provide a controlled environment in which the effects of different initial conditions, boundary conditions and buoyancy sources on core dynamics and the resulting magnetic field can be investigated.

Most numerical simulations solve the non-dimensional governing equations within a spherical shell of thickness $L = r_o - r_i$, where $r_o$ and $r_i$ are the outer and inner radii, respectively. The shell rotates about the vertical axis $\hat{\boldsymbol{z}}$ with angular frequency $\Omega$. The Boussinesq approximation is adopted, under which density variations are neglected except in terms multiplied by the acceleration due to gravity. The fluid is also assumed to be incompressible, with constant material properties. The non-dimensional governing equations are therefore:

- The Navier–Stokes equation, which balances fluid acceleration with the forces acting on the fluid

$$\underbrace{\frac{Ek}{Pm}\left[\frac{\partial \boldsymbol{u}^*}{\partial t} + (\boldsymbol{u}^* \cdot \nabla)\boldsymbol{u}^*\right]}_{\text{Inertia}} + \underbrace{\hat{\boldsymbol{z}} \times \boldsymbol{u}^*}_{\text{Coriolis}} = \underbrace{-\nabla P^*}_{\text{Pressure}} + \underbrace{\frac{Pm}{Pr} RaC^* \boldsymbol{r}}_{\text{Buoyancy}} + \underbrace{(\nabla \times \boldsymbol{B}^*) \times \boldsymbol{B}^*}_{\text{Lorentz}} + \underbrace{Ek\nabla^2 \boldsymbol{u}^*}_{\text{Viscosity}} \tag{1.35}$$

- The induction equation, which describes magnetic-field generation by fluid motion and its decay by diffusion

$$\frac{\partial \boldsymbol{B}^*}{\partial t} = \underbrace{\nabla \times (\boldsymbol{u}^* \times \boldsymbol{B}^*)}_{\text{Advection}} + \underbrace{\nabla^2 \boldsymbol{B}^*}_{\text{Diffusion}} \tag{1.36}$$

- The codensity equation, which describes changes in the density anomaly $C^*$ through advection, diffusion and sources or sinks $S$, such as radiogenic heating

$$\frac{\partial C^*}{\partial t} = -\boldsymbol{u}^* \cdot \nabla C^* + \frac{Pm}{Pr}\nabla^2 C^* + S \tag{1.37}$$

- The continuity equations, which describe the incompressibility of the core fluid and the absence of magnetic monopoles

$$\nabla \cdot \boldsymbol{u}^* = \nabla \cdot \boldsymbol{B}^* = 0 \tag{1.38}$$

where $t$ is time, $\boldsymbol{r}$ is the radial position vector within the spherical coordinate system, and $P^*$ is the modified pressure perturbation.

To compare the simulation outputs with observations, the dimensionless magnetic field and time must be converted into dimensional quantities. For the scaling used in Equations (1.35)–(1.38), the dimensional magnetic field may be written as

$$\boldsymbol{B} = \sqrt{\rho \mu_0 \Omega \eta}\, \boldsymbol{B}^*, \tag{1.39}$$

where $\rho$ is the fluid density, $\mu_0$ is the permeability of free space, $\Omega$ is the rotation rate and $\eta$ is the magnetic diffusivity. The dimensional time may be obtained using either the magnetic-diffusion timescale, $\tau_d = L^2/\eta$, or the advection timescale, $\tau_a = L/U$, where $L$ is the characteristic length scale, and $U$ is the characteristic velocity scale. Since the magnetic Reynolds number satisfies $R_m = \tau_d/\tau_a$, these scalings are equivalent only when the simulation and the Earth have comparable values of $R_m$. A choice of timescale is therefore required when comparing the temporal evolution of a simulation with observations of the Earth's magnetic field.

The relative importance of the forces and processes represented in these equations is described by a series of dimensionless parameters. Their definitions, together with typical values for the Earth's core, are given in Table 1.2. Numerical simulations cannot currently reach Earth-like values of these dimensionless control parameters because of the wide range of spatial and temporal scales involved. Resolving these scales is computationally expensive, particularly when simulations must be run for long enough to capture the long-term behaviour of the magnetic field. Numerical simulations can reproduce the Earth's Prandtl number but cannot reach Earth-like Ekman and magnetic Prandtl numbers. Enhanced diffusion at small scales reduces the computational cost and allows long-term behaviour to be studied at lower Ekman numbers, although the approach requires careful validation (Aubert et al., 2017; Frasson et al., 2025). Shorter simulations can reach more Earth-like parameter regimes but cannot capture longer-term behaviour such as excursions and reversals (Aubert, 2023).

| **Dimensionless Parameter** | **Definition** | **Earth Value** |
|---|---|---|
| Prandtl, $Pr$ | Ratio of momentum diffusivity to thermal diffusivity | $\sim 1$ |
| Magnetic Prandtl, $Pm$ | Ratio of momentum diffusivity to magnetic diffusivity | $\sim 10^{-6}$ |
| Ekman, $Ek$ | Ratio of viscous to Coriolis forces | $10^{-15}$ |
| Rayleigh, $Ra$ | Strength of the buoyancy source driving convection | Larger than $Ra_c$ needed for onset of convection |

Table 1.2: Summary of dimensionless control parameters used in Equations (1.35)-(1.38). Value for the Earth taken from the review article by Wicht and Sanchez (2019).

The force balances in the Earth's core can also be considered in these simulations. One possible large-scale dynamical regime is the quasi-geostrophic magneto-Archimedes-Coriolis (QG–MAC) balance. The Coriolis force is largely balanced by the pressure gradient, while the remaining Coriolis force is balanced by the Lorentz and buoyancy forces (Schwaiger et al., 2019; Nakagawa and Davies, 2022). Aubert et al. (2017) and Aubert (2019) showed that this balance persists along a path through parameter space towards the conditions expected in Earth's core, with inertial and viscous effects becoming progressively less important. Aubert (2026) subsequently reached 100% of this path, confirming that the QG–MAC balance is maintained when the core density, rotation rate, convective power and electrical conductivity are matched, while viscous losses remain minor. Although geodynamo simulations cannot access the Earth's true parameter regime, some reproduce aspects of the observed field morphology and variability (Rogers et al., 2026).

## 1.11 Core–Mantle Boundary Coupling

Seismic tomography provides the principal observational constraint on the structure of the lowermost mantle. It uses differences between the observed travel times of seismic waves and those predicted by a reference Earth model, such as the Preliminary Reference Earth Model (PREM) (Dziewonski and Anderson, 1981), to map variations in seismic-wave velocity and attenuation (Anderson and Dziewonski, 1984; Rawlinson et al., 2010), which can be related to structures in the mantle with different temperature or composition. A schematic of interpretations of these structures are shown in Figure 1.16, and include the two large low-velocity provinces (LLVPs) beneath Africa and the Pacific, smaller ultra-low-velocity zones (ULVZs) immediately above the boundary, the D″ discontinuity and regions associated with subducted material (Dziewonski et al., 1977; Masters et al., 1996; Koelemeijer, 2021; Ritsema and Lekić, 2020).

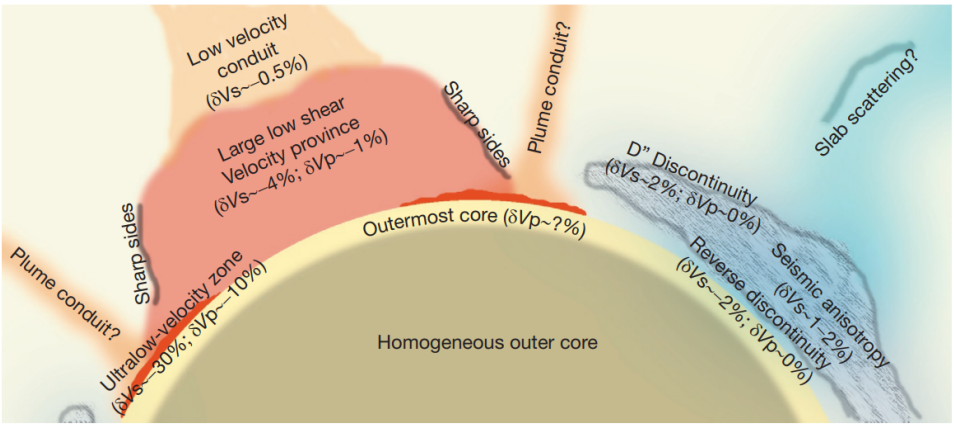



Figure 1.16: Schematic of the principal structures identified by seismic observations in the core–mantle boundary region, taken from Hernlund and McNamara (2015).

### 1.11.1 Large Low-Velocity Provinces (LLVPs)

The most prominent large-scale structures in the lowermost mantle are the large low-velocity provinces (LLVPs), two approximately antipodal regions beneath Africa (named Jason) and the Pacific (named Tuzo) that cover around 25% of the core surface and are characterised by reduced shear- and compressional-wave velocities (Masters et al., 1996; Ritsema and Lekić, 2020; Koelemeijer, 2021). Although both extend for thousands of kilometres laterally, the African LLVP is approximately 30% more voluminous and rises around 1000 km higher above the CMB than the Pacific LLVP (Doucet et al., 2020; Yuan and Li, 2022). Geodynamic modelling suggests that this greater height may reflect a lower intrinsic density and reduced stability, indicating

differences in the composition and evolution of the two structures (Yuan and Li, 2022).

The physical origin of the LLVPs remains uncertain. Their reduced seismic velocities may reflect elevated temperature with respect to the ambient mantle, but their sharp lateral boundaries suggest that they are compositionally distinct, rather than being purely thermal anomalies (Hernlund and McNamara, 2015; Mosca et al., 2012; McNamara, 2019). The differences in their morphology and composition further suggest that the two LLVPs have experienced distinct interactions with the surrounding mantle and followed different evolutionary histories (Doucet et al., 2020; Yuan and Li, 2022). Proposed origins include the accumulation of subducted material, remnants of an early magma ocean and long-lived thermochemical structures formed during mantle evolution (Hernlund and McNamara, 2015; McNamara, 2019). However, the relative contributions of thermal and compositional variations remain difficult to determine from seismic observations alone.

### 1.11.2 Outermost-Core Stratification

Outermost-core stratification refers to a gravitationally stable region at the top of the outer core in which coherent radial convective motions from the deeper core are inhibited from reaching the CMB. This stability may result from the accumulation of light elements in the outermost core, a subadiabatic thermal gradient, or a combination of both (Fearn and Loper, 1981; Labrosse et al., 1997; Buffett and Seagle, 2010). The evidence for such a layer is mixed, with seismic studies arguing both for (Lay and Young, 1990; Helffrich and Kaneshima, 2010; Kaneshima, 2018; Ma and Tkalčić, 2023) and against (Alexandrakis and Eaton, 2010; Irving et al., 2018; Munch et al., 2026) the existence of a globally stratified layer. Geomagnetic evidence is also inconclusive. Early analyses found that the observed secular variation was consistent with negligible horizontal flow divergence and therefore with stable stratification (Whaler, 1980). More recent studies have identified possible signatures of radial motion and found that a large-scale purely toroidal flow cannot reproduce the observed field (Gubbins, 2007; Amit, 2014; Lesur et al., 2015). However, these results remain consistent with weaker stratification that permits some poloidal motion and are sensitive to the contribution of magnetic diffusion. Stratification may also occur regionally rather than as a continuous global layer. In geodynamo simulations, lateral variations in CMB heat flux can produce stable regional inversion layers beneath anomalously hot regions of the lowermost mantle, particularly beneath the African and Pacific LLVPs, while the remainder of the outer core continues to convect (Mound et al., 2019; Mound and Davies, 2020; Naskar et al.,

2026). These layers may extend several hundred kilometres into the core and, when horizontally averaged, produce a radial temperature profile resembling global stratification (Mound et al., 2019; Mound and Davies, 2020; Deschamps et al., 2026; Naskar et al., 2026).

### 1.11.3 Heat Flux Heterogeneity

Thermal coupling across the CMB provides a mechanism through which lowermost-mantle structure may influence core flow and, consequently, the geomagnetic field. Lateral temperature variations in the lowermost mantle produce a heterogeneous CMB heat flux. Anomalously hot mantle reduces the temperature difference across the boundary and therefore the local heat flux out of the core, whereas colder regions associated with mantle downwellings or remnants of subducted slabs produce regional heat-flux maxima (Mound and Davies, 2023; Dannberg et al., 2024). If the reduced seismic velocities within the LLVPs reflect elevated temperatures, their footprints would therefore be associated with reduced CMB heat flux. CMB heat-flux patterns are commonly estimated from seismic tomography by relating seismic-velocity anomalies to temperature, as in models derived from Masters et al. (1996). However, seismic velocity is also affected by composition. The difficulty of separating these effects means that the amplitude and detailed structure of CMB heat-flux variations remain uncertain (Zhao et al., 2015; Mosca et al., 2012; Mound and Davies, 2023). Mantle-convection simulations provide an alternative approach by calculating heat flux from the modelled thermal structure, although the results depend on the assumed mantle properties and plate reconstruction.

Mantle structure evolves slowly relative to the outer core, so CMB heat-flux heterogeneity provides a persistent, but gradually changing, thermal boundary condition on the geodynamo (Mound and Davies, 2023; Dannberg et al., 2024). True polar wander may redistribute this pattern relative to the rotation axis, altering the balance between polar and equatorial cooling (Frasson et al., 2024). Numerical simulations show that heterogeneous heat flux can reorganise core convection and has been linked to persistent magnetic flux patches, secular variation, non-zonal field structure and changes in reversal frequency (Bloxham and Gubbins, 1987; Amit et al., 2010; Glatzmaier et al., 1999; Frasson et al., 2025; Mound and Davies, 2023). Where the local heat flux is sufficiently low, convection in the uppermost core may also be reduced, providing a mechanism for the regional stratification discussed above (Mound et al., 2019; Mound and Davies, 2020; Naskar et al., 2026). Evidence for a persistent mantle influence on the geomagnetic field is nevertheless inconclusive. Long-term field models and archeomagnetic

observations indicate that magnetic flux patches can migrate and are not permanently fixed at preferred longitudes (Amit et al., 2011; Panovska et al., 2019). Similar magnetic behaviour may also arise from different dynamo regimes or boundary conditions, making it difficult to isolate the contribution of mantle influence from the magnetic field alone.

## References

Alexandrakis, C., Eaton, D.W., 2010. Precise seismic-wave velocity atop earth's core: No evidence for outer-core stratification. Physics of the Earth and Planetary Interiors 180, 59–65. URL: https://www.sciencedirect.com/science/article/pii/S0031920110000452, doi:https://doi.org/10.1016/j.pepi.2010.02.011.

Alken, P., Chulliat, A., Nair, M., 2021a. Noaa/ncei and university of colorado candidate models for igrf-13. Earth, Planets and Space 73, 44. URL: https://doi.org/10.1186/s40623-020-01313-z, doi:10.1186/s40623-020-01313-z.

Alken, P., Thébault, E., Beggan, C.D., Amit, H., Aubert, J., Baerenzung, J., Bondar, T.N., Brown, W.J., Califf, S., Chambodut, A., Chulliat, A., Cox, G.A., Finlay, C.C., Fournier, A., Gillet, N., Grayver, A., Hammer, M.D., Holschneider, M., Huder, L., Hulot, G., Jager, T., Kloss, C., Korte, M., Kuang, W., Kuvshinov, A., Langlais, B., Léger, J.M., Lesur, V., Livermore, P.W., Lowes, F.J., Macmillan, S., Magnes, W., Mandea, M., Marsal, S., Matzka, J., Metman, M.C., Minami, T., Morschhauser, A., Mound, J.E., Nair, M., Nakano, S., Olsen, N., Pavón-Carrasco, F.J., Petrov, V.G., Ropp, G., Rother, M., Sabaka, T.J., Sanchez, S., Saturnino, D., Schnepf, N.R., Shen, X., Stolle, C., Tangborn, A., Tøffner-Clausen, L., Toh, H., Torta, J.M., Varner, J., Vervelidou, F., Vigneron, P., Wardinski, I., Wicht, J., Woods, A., Yang, Y., Zeren, Z., Zhou, B., 2021b. International geomagnetic reference field: the thirteenth generation. Earth, Planets and Space 73, 49. doi:10.1186/s40623-020-01288-x.

Amit, H., 2014. Can downwelling at the top of the earth's core be detected in the geomagnetic secular variation? Physics of the Earth and Planetary Interiors 229, 110–121. doi:10.1016/j.pepi.2014.01.012.

Amit, H., Aubert, J., Hulot, G., 2010. Stationary, oscillating or drifting mantle-driven geomagnetic flux patches? Journal of Geophysical Research: Solid Earth 115. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2009JB006542, doi:https://doi.org/10.1029/2009JB006542, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2009JB006542.

Amit, H., Christensen, U.R., 2008. Accounting for magnetic diffusion in core flow inversions from geomagnetic secular variation. Geophysical Journal International 175, 913–924. URL: https://doi.org/10.1111/j.1365-246X.2008.03948.x, doi:10.1111/j.1365-246X.2008.03948.x, arXiv:https://academic.oup.com/gji/article-pdf/175/3/913/5997300/175-3-913.pdf.

Amit, H., Korte, M., Aubert, J., Constable, C., Hulot, G., 2011. The time-dependence of intense archeomagnetic flux patches. Journal of Geophysical Research: Solid Earth 116. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011JB008538, doi:https://doi.org/10.1029/2011JB008538, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2011JB008538.

Amit, H., Olson, P., 2004. Helical core flow from geomagnetic secular variation. Physics of the Earth and Planetary Interiors 147, 1–25. doi:10.1016/j.pepi.2004.02.006.

Amit, H., Pais, M.A., 2013. Differences between tangential geostrophy and columnar flow. Geophysical Journal International 194, 145–157. doi:10.1093/gji/ggt077.

Anderson, D.L., Dziewonski, A.M., 1984. Seismic tomography. Scientific American 251, 60–71. URL: http://www.jstor.org/stable/24969456.

Aubert, J., 2019. Approaching Earth's core conditions in high-resolution geodynamo simulations. Geophysical Journal International 219, S137–S151. URL: https://doi.org/10.1093/gji/ggz232, doi:10.1093/gji/ggz232.

Aubert, J., 2023. State and evolution of the geodynamo from numerical models reaching the physical conditions of earth's core. Geophysical Journal International 235,

468–487. URL: https://doi.org/10.1093/gji/ggad229, doi:10.1093/gji/ggad229, arXiv:https://academic.oup.com/gji/article-pdf/235/1/468/50615243/ggad229.pdf.

Aubert, J., 2026. Geodynamo simulations spanning millennia in the physical conditions of earth's core. Journal of Studies of Earth's Deep Interior Volume 2. URL: http://jsedi.episciences.org/17790, doi:10.46298/jsedi.17790.

Aubert, J., Gastine, T., Fournier, A., 2017. Spherical convective dynamos in the rapidly rotating asymptotic regime. Journal of Fluid Mechanics 813, 558–593. doi:10.1017/jfm.2016.789.

Backus, 1968. Kinematics of geomagnetic secular variation in a perfectly conducting core. Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences 263, 239–266.

Backus, Le Mouël, J., 1986. The region on the core—mantle boundary where a geostrophic velocity field can be determined from frozen-flux magnetic data. Geophysical Journal International 85, 617–628. doi:10.1111/j.1365-246X.1986.tb04536.x.

Baerenzung, J., Holschneider, M., Wicht, J., Lesur, V., Sanchez, S., 2020. The kalmag model as a candidate for igrf-13. Earth, Planets and Space 72, 163. URL: https://doi.org/10.1186/s40623-020-01295-y, doi:10.1186/s40623-020-01295-y.

Baratchart, L., Gerhards, C., 2017. Comments on uniqueness in internal–external separation. Inverse Problems .

Baumjohann, W., Nakamura, R., 2007. Magnetospheric contributions to the terrestrial magnetic field, in: Treatise on Geophysics. Elsevier, Amsterdam. volume 5, pp. 77–92. doi:10.1016/B978-044452748-6.00088-2.

Beggan, C., Whaler, K., 2008. Core flow modelling assumptions. Physics of the Earth and Planetary Interiors 167, 217–222. URL: https://www.sciencedirect.com/science/article/pii/S0031920108000812, doi:https://doi.org/10.1016/j.pepi.2008.04.011.

Beggan, C.D., Kloss, C., Amblard, P., Amezza, T., Amit, H., Aubert, J., Baerenzung, J., Bregnhøj Lauridsen, J., Brown, W., Califf, S., Campuzano, S.A., Chauvet, L., Chulliat, A., Claveau, R., Coïsson, P., da Silva, M., Davies, C., Deborde, R., Dobrovolsky, M., Fillion, M., Finlay, C.C., Firsov, I., Fournier, A., Gailler, L., Gillet, N., Gomez Perez, N., Grayver, A., Gwirtz, K., Hamoudi, M., Holme, R., Holschneider, M., Hulot, G., Jager, T., Jiang, Y., Kada, M., Korte, M., Krasnoperov, R., Kuang, W., Kudin, D., Léger, J.M., Lemgharb, A., Lesur, V., Liu, P., Livermore, P.W., Madsen, F.D., Mandea, M., Matsushima, M., Matzka, J., Michaelis, I., Minami, T., Nair, M., Nakano, S., Oehler, J.F., Olsen, N., Pavón-Carrasco, F.J., Presnyakov, S., Qian, B., Rang, X., Rogers, H.F., Rother, M., Rouxel, D., Sabaka, T., Sato, S., Saynisch-Wagner, J., Schanner, M., Serrano, M., Shakespeare-Rees, N., Shi, Y., Smith, A., Soloviev, A., Terra-Nova, F., Thébault, E., Tøffner-Clausen, L., Toh, H., Wang, J., Wardinski, I., Watson, C., Whaler, K., Xiong, C., Yan, Q., Yang, Y., Yao, H., Zeren, Z., Zhang, K., Zhou, B., 2026. International geomagnetic reference field: The fourteenth generation. Earth, Planets and Space 78, 127. URL: https://doi.org/10.1186/s40623-025-02360-0, doi:10.1186/s40623-025-02360-0.

Beggan, C.D., Saarimäki, J., Whaler, K.A., Simons, F.J., 2013. Spectral and spatial decomposition of lithospheric magnetic field models using spherical slepian functions. Geophysical Journal International 193, 136–148. doi:10.1093/gji/ggs122.

Bloxham, Jackson, 1991. Fluid flow near the surface of earth's outer core. Reviews of Geophysics 29, 97–120. doi:10.1029/90RG02470.

Bloxham, J., Gubbins, D., 1987. Thermal core– mantle interactions. Nature 325, 511–513. URL: https://doi.org/10.1038/325511a0, doi:10.1038/325511a0.

Bloxham, J., Jackson, A., 1992. Time-dependent mapping of the magnetic field at the core-mantle boundary. Journal of Geophysical Research: Solid Earth 97, 19537–

19563. URL: https://onlinelibrary.wiley.com/doi/abs/10.1029/92JB01591, doi:10.1029/92JB01591. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/92JB01591.

Brown, W., Beggan, C., Gomez Perez, N., Hammer, M., Finlay, C., Hulot, G., 2022. Geomagnetic virtual observatories for investigating sub-annual core field variation. URL: https://lps22.esa.int/frontend/index.php?page_id=18477.

Brown, W.J., Mound, J.E., Livermore, P.W., 2013. Jerks abound: An analysis of geomagnetic observatory data from 1957 to 2008. Physics of the Earth and Planetary Interiors 223, 62–76. URL: https://www.sciencedirect.com/science/article/pii/S0031920113000708, doi:10.1016/j.pepi.2013.06.001.

Buffett, B.A., Seagle, C.T., 2010. Stratification of the top of the core due to chemical interactions with the mantle. Journal of Geophysical Research: Solid Earth 115, B04407. URL: https://doi.org/10.1029/2009JB006751, doi:10.1029/2009JB006751.

Bullard, E.C., 1950. The transfer of heat from the core of the earth. Geophysical Journal International 6, 36–41. doi:10.1111/j.1365-246X.1950.tb02971.x.

Chambodut, A., et al., 2005. Wavelet frames: An alternative to spherical harmonic representation of magnetic fields. Geophysical Journal International 163, 875–899.

Colomb, F., Varotto, C., 2003. Sac-c and the am constellation: three years of achievements, in: International Conference on Recent Advances in Space Technologies, 2003. RAST '03. Proceedings of, pp. 119–124. doi:10.1109/RAST.2003.1303893.

Constable, C., 2019. Temporal field variations: from reversals to very-low-frequency waves. Cambridge University Press. pp. 181–185.

Constable, C., Constable, S., 2023. A grand spectrum of the geomagnetic field. Physics of the Earth and Planetary Interiors 344, 107090. URL: https://www.sciencedirect.com/science/article/pii/S0031920123001164, doi:https://doi.org/10.1016/j.pepi.2023.107090.

Constable, C., Korte, M., 2015. Centennial- to millennial-scale geomagnetic field variations, in: Treatise on Geophysics. 2 ed.. Elsevier.

Constable, C.G., Constable, S.C., 2004. Satellite magnetic field measurements: applications in studying the deep earth. The state of the planet: Frontiers and challenges in geophysics 150, 147–159.

Cormier, V.F., Bergman, M.I., Olson, P.L., 2022. Chapter 3 - Geodynamo and geomagnetic basics. Elsevier. p. 75–113. URL: https://www.sciencedirect.com/science/article/pii/B9780128114001000045, doi:10.1016/B978-0-12-811400-1.00004-5.

Dannberg, J., Gassmöller, R., Thallner, D., LaCombe, F., Sprain, C., 2024. Changes in core–mantle boundary heat flux patterns throughout the supercontinent cycle. Geophysical Journal International 237, 1251–1274. doi:10.1093/gji/ggae075.

Davies, C., Constable, C., 2017. Geomagnetic spikes on the core–mantle boundary. Nature Communications 8, 15593. doi:10.1038/ncomms15593.

Deconinck, F., Hulot, G., Léger, J.M., Clausen, L.B.N., Pastena, M., Lejault, J.P., Wilkinson, J., Caballero, D.M., Tostado, M.V., Coïsson, P., Jager, T., Nieto, P., Jørgensen, J.L., 2025. NanoMagSat status: a 3x16U low-Earth orbit constellation to monitor the Earth magnetic field and the ionospheric environment, in: Petrozzi-Ilstad, M. (Ed.), Small Satellites Systems and Services Symposium (4S 2024), International Society for Optics and Photonics. SPIE. p. 135462K. URL: https://doi.org/10.1117/12.3062471, doi:10.1117/12.3062471.

Deguen, Alboussière, T., Cardin, P., 2013. Thermal convection in earth's inner core with phase change at its boundary. Geophysical Journal International 194, 1310–1334. doi:10.1093/gji/ggt202.

Deschamps, F., Guerrero, J.M., Amit, H., Terra-Nova, F., Hsieh, W.P., 2026. Neg-

ative core–mantle boundary heat flux beneath low-shear-wave-velocity provinces. Nature Geoscience URL: https://doi.org/10.1038/s41561-026-02018-w, doi:10.1038/s41561-026-02018-w.

Doucet, L.S., Li, Z.X., Gamal El Dien, H., Pourteau, A., Murphy, J.B., Collins, W.J., Mattielli, N., Olierook, H.K.H., Spencer, C.J., Mitchell, R.N., 2020. Distinct formation history for deep-mantle domains reflected in geochemical differences. Nature Geoscience 13, 511–515. URL: https://doi.org/10.1038/s41561-020-0599-9, doi:10.1038/s41561-020-0599-9.

Dziewonski, A.M., Anderson, D.L., 1981. Preliminary reference earth model. Physics of the Earth and Planetary Interiors 25, 297–356. URL: https://www.sciencedirect.com/science/article/pii/0031920181900467, doi:https://doi.org/10.1016/0031-9201(81)90046-7.

Dziewonski, A.M., Hager, B.H., O'Connell, R.J., 1977. Large-scale heterogeneities in the lower mantle. Journal of Geophysical Research (1896-1977) 82, 239–255. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JB082i002p00239, doi:https://doi.org/10.1029/JB082i002p00239, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/JB082i002p00239.

Fearn, D.R., Loper, D.E., 1981. Compositional convection and stratification of Earth's core. Nature 289, 393–394. URL: https://doi.org/10.1038/289393a0, doi:10.1038/289393a0.

Finlay, C.C., Aubert, J., Gillet, N., 2016. Gyre-driven decay of the Earth's magnetic dipole. Nature Communications 7, 10422. URL: https://www.nature.com/articles/ncomms10422, doi:10.1038/ncomms10422. number: 1 Publisher: Nature Publishing Group.

Finlay, C.C., Gillet, N., Aubert, J., Livermore, P.W., Jault, D., 2023. Gyres, jets and waves in the earth's core. Nature Reviews Earth & Environment 4, 377–392. URL: https://www.nature.com/articles/s43017-023-00425-w, doi:10.1038/s43017-023-00425-w.

Finlay, C.C., Kloss, C., Olsen, N., Hammer, M.D., Tøffner-Clausen, L., Grayver, A., Kuvshinov, A., 2020. The chaos-7 geomagnetic field model and observed changes in the south atlantic anomaly. Earth, Planets and Space 72, 156. URL: https://doi.org/10.1186/s40623-020-01252-9, doi:10.1186/s40623-020-01252-9.

Finlay, C.C., Lesur, V., Thébault, E., Vervelidou, F., Morschhauser, A., Shore, R., 2017. Challenges handling magnetospheric and ionospheric signals in internal geomagnetic field modelling. Space Science Reviews 206, 157–189. doi:10.1007/s11214-016-0285-9.

Fournier, A., Aubert, J., Lesur, V., Ropp, G., 2021. A secular variation candidate model for igrf-13 based on swarm data and ensemble inverse geodynamo modelling. Earth, Planets and Space 73, 43. URL: https://doi.org/10.1186/s40623-020-01309-9, doi:10.1186/s40623-020-01309-9.

Fournier, A., Hulot, G., Jault, D., Kuang, W., Tangborn, A., Gillet, N., Canet, E., Aubert, J., Lhuillier, F., 2010. An introduction to data assimilation and predictability in geomagnetism. Space Science Reviews 155, 247–291. URL: https://doi.org/10.1007/s11214-010-9669-4, doi:10.1007/s11214-010-9669-4.

Frasson, T., Labrosse, S., Nataf, H.C., Coltice, N., Flament, N., 2024. On the impact of true polar wander on heat flux patterns at the core–mantle boundary. Solid Earth 15, 617–637. URL: https://se.copernicus.org/articles/15/617/2024/, doi:10.5194/se-15-617-2024.

Frasson, T., Schaeffer, N., Nataf, H.C., Labrosse, S., 2025. Geomagnetic dipole stability and zonal flows controlled by mantle heat flux heterogeneities. Geophysical Journal International 240, 1481–1504. URL: https://doi.org/10.1093/gji/ggae457, doi:10.1093/gji/ggae457, arXiv:https://academic.oup.com/gji/article-pdf/240/3/1481/61290993/ggae457.pdf.

Friis-Christensen, E., Lühr, H., Hulot, G., 2006. Swarm: A constellation to study the earth's

magnetic field. Earth, Planets and Space 58, 351–358. doi:10.1186/BF03351933.

Gao, Y., Wang, Z.T., Livermore, P.W., Rogers, H.F., Liu, C., 2025. Mscm: A geomagnetic model derived from swarm, cses, and mss-1 satellite data and the evolution of the south atlantic anomaly. Earth and Planetary Physics 9, 564–576. doi:10.26464/epp2025030.

Gauss, C.F., 1839. Allgemeine theorie des erdmagnetismus, in: Resultate aus den Beobachtungen des Magnetischen Vereins im Jahre 1838. Weidmannsche Buchhandlung, Leipzig.

Gilbert, W., 1600. De Magnete, Magneticisque Corporibus, et de Magno Magnete Tellure. Peter Short, London.

Gillet, Gerick, F., Jault, D., Schwaiger, T., Aubert, J., Istas, M., 2022. Satellite magnetic data reveal interannual waves in earth's core. Proceedings of the National Academy of Sciences of the United States of America 119, e2115258119. doi:10.1073/pnas.2115258119.

Gillet, Schaeffer, N., Jault, D., 2012. Rationale and geophysical evidence for quasi-geostrophic rapid dynamics within the earth's outer core. Physics of the Earth and Planetary Interiors 202–203, 78–88. doi:10.1016/j.pepi.2012.03.006.

Gillet, N., Barrois, O., Finlay, C.C., 2015. Stochastic forecasting of the geomagnetic field from the cov-obs.x1 geomagnetic field model, and candidate models for igrf-12. Earth, Planets and Space 67, 71. URL: https://doi.org/10.1186/s40623-015-0225-z, doi:10.1186/s40623-015-0225-z.

Gillet, N., Dall'Asta, F., Amblard, P.O., Claveau, R., Aubert, J., 2024. Waves in earth's core and geomagnetic field forecast. Physics of the Earth and Planetary Interiors 357, 107284. URL: https://www.sciencedirect.com/science/article/pii/S0031920124001420, doi:https://doi.org/10.1016/j.pepi.2024.107284.

Glaßmeier, K.H., Soffel, H., Negendank, J.F.W., 2009. Geomagnetic Field Variations. Advances in Geophysical and Environmental Mechanics and Mathematics. 1 ed., Springer Berlin, Heidelberg. URL: https://doi.org/10.1007/978-3-540-76939-2, doi:10.1007/978-3-540-76939-2.

Glatzmaier, G.A., Coe, R.S., Hongre, L., Roberts, P.H., 1999. The role of the Earth's mantle in controlling the frequency of geomagnetic reversals. Nature 401, 885–890. doi:10.1038/44776.

Griffiths, D.J., 2017. Introduction to Electrodynamics. 4 ed., Cambridge University Press. doi:10.1017/9781108333511.

Grüne, C.R., Whaler, K.A., Madsen, F.D., 2025. Detecting low-latitude outer core-surface waves with 25 years of satellite secular variation data. Physics of the Earth and Planetary Interiors 367, 107435. URL: https://www.sciencedirect.com/science/article/pii/S0031920125001293, doi:https://doi.org/10.1016/j.pepi.2025.107435.

Gubbins, 2007. Geomagnetic constraints on stratification at the top of earth's core. Earth, Planets and Space 59, 661–664. doi:10.1186/BF03352728.

Gubbins, D., 1996. A formalism for the inversion of geomagnetic data for core motions with diffusion. Physics of the Earth and Planetary Interiors 98, 193–206. doi:10.1016/S0031-9201(96)03187-1.

Gubbins, D., Bloxham, J., 1985. Geomagnetic field analysis—iii. magnetic fields on the core–mantle boundary. Geophysical Journal International 80, 695–713. doi:10.1111/j.1365-246X.1985.tb05119.x.

Gubbins, D., Herrero-Bervera, E. (Eds.), 2007. Encyclopedia of geomagnetism and paleomagnetism. Encyclopedia of earth sciences series, Springer, Dordrecht. OCLC: ocn159554151.

Halley, E., 1701. A new and correct chart shewing the variations of the compass. Map of magnetic declination.

Hammer, M.D., Cox, G.A., Brown, W.J., Beggan, C.D., Finlay, C.C., 2021. Geomagnetic

virtual observatories: monitoring geomagnetic secular variation with the swarm satellites. Earth, Planets and Space 73, 54. doi:10.1186/s40623-021-01357-9.

Hammer, M.D., Finlay, C.C., 2018. Local averages of the core–mantle boundary magnetic field from satellite observations. Geophysical Journal International 216, 1901–1918. URL: https://doi.org/10.1093/gji/ggy515, doi:10.1093/gji/ggy515, arXiv:https://academic.oup.com/gji/article-pdf/216/3/1901/27493010/ggy515.pdf.

Helffrich, G., Kaneshima, S., 2010. Outer-core compositional stratification from observed core wave speed profiles. Nature 468, 807–810. URL: https://doi.org/10.1038/nature09636, doi:10.1038/nature09636.

Hernlund, J.W., McNamara, A.K., 2015. The core–mantle boundary region, in: Treatise on Geophysics. 2 ed.. Elsevier. volume 7, pp. 461–519. doi:10.1016/B978-0-444-53802-4.00136-6.

Hirose, K., Wood, B., Vočadlo, L., 2021. Light elements in the Earth's core. Nature Reviews Earth & Environment 2, 645–658. URL: https://doi.org/10.1038/s43017-021-00203-6, doi:10.1038/s43017-021-00203-6.

Hollerbach, Gubbins, 2007. Inner Core Tangent Cylinder. Springer Netherlands, Dordrecht. p. 430–433. URL: https://doi.org/10.1007/978-1-4020-4423-6_154, doi:10.1007/978-1-4020-4423-6_154.

Hollerbach, R., 1996. On the theory of the geodynamo. Physics of the Earth and Planetary Interiors 98, 163–185. doi:10.1016/S0031-9201(96)03185-8.

Holme, R., 2015. 8.04 - large-scale flow in the core, in: Schubert, G. (Ed.), Treatise on Geophysics (Second Edition). second edition ed.. Elsevier, Oxford, pp. 91–113. URL: https://www.sciencedirect.com/science/article/pii/B978044453802400138X, doi:https://doi.org/10.1016/B978-0-444-53802-4.00138-X.

Holme, R., Olsen, N., 2006. Core surface flow modelling from high-resolution secular variation. Geophysical Journal International 166, 518–528. URL: https://doi.org/10.1111/j.1365-246X.2006.03033.x, doi:10.1111/j.1365-246X.2006.03033.x, arXiv:https://academic.oup.com/gji/article-pdf/166/2/518/1499206/166-2-518.pdf.

Holschneider, M., et al., 2016. Bayesian approach to separating internal and external geomagnetic fields. Journal of Geophysical Research .

Huang, X., Gerhards, C., Ren, Z., 2026. Global versus local internal–external field separation on the sphere: a hardy–hodge perspective. arXiv preprint arXiv:2603.04886 .

Huder, L., Gillet, N., Finlay, C.C., Hammer, M.D., Tchoungui, H., 2020. Cov-obs.x2: 180 years of geomagnetic field evolution from ground-based and satellite observations. Earth, Planets and Space 72, 160. URL: https://doi.org/10.1186/s40623-020-01194-2, doi:10.1186/s40623-020-01194-2.

Hulot, G., Sabaka, T.J., Olsen, N., Fournier, A., 2015. The present and future geomagnetic field, in: Treatise on Geophysics. 2 ed.. Elsevier. volume 5, pp. 33–78. doi:10.1016/B978-0-444-53802-4.00096-8.

INTERMAGNET, . Intermagnet metadata map. https://intermagnet.org/metadata/map. Accessed: 2026-06-04.

Irving, J.C.E., Cottaar, S., Lekić, V., 2018. Seismically determined elastic parameters for earth's outer core. Science Advances 4, eaar2538. URL: https://www.science.org/doi/abs/10.1126/sciadv.aar2538, doi:10.1126/sciadv.aar2538, arXiv:https://www.science.org/doi/pdf/10.1126/sciadv.aar2538.

Jackson, A., Finlay, C.C., 2015. Geomagnetic secular variation and its applications to the core, in: Schubert, G. (Ed.), Treatise on Geophysics. 2 ed.. Elsevier, Oxford. volume 5, pp. 137–184. doi:10.1016/B978-0-444-53802-4.00099-3.

Jackson, A., Jonkers, A.R.T., Walker, M.R., 2000. Four centuries of geomagnetic secular variation from historical records. Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 358, 957–990. doi:10.1098/rsta.2000.0569.

Jackson, J.E., Vette, J.I., 1975. OGO Program Summary. Technical Report NASA-SP-7601. NASA Goddard Space Flight Center. Greenbelt, MD. NASA Special Publication.

Jacobs, J.A., 1994. Reversals of the Earth's Magnetic Field. Cambridge University Press.

Jiang, Y., Finlay, C.C., Olsen, N., Tøffner-Clausen, L., Yan, Q., Zhang, K., 2024. Macau scientific satellite-1 initial magnetic field model. Geophysical Research Letters 51, e2024GL112305. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2024GL112305, doi:https://doi.org/10.1029/2024GL112305, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2024GL112305. e2024GL112305 2024GL112305.

Johnson, C.L., McFadden, P.L., 2015. The time-averaged field and paleosecular variation, in: Treatise on Geophysics. 2 ed.. Elsevier.

Jonkers, A.R.T., 2003. Earth's Magnetism in the Age of Sail. Johns Hopkins University Press.

Kaneshima, S., 2018. Array analyses of smks waves and the stratification of earth's outermost core. Physics of the Earth and Planetary Interiors 276, 234–246. doi:10.1016/j.pepi.2017.03.006.

Kauristie, K., Morschhauser, A., Olsen, N., Finlay, C.C., McPherron, R.L., Gjerloev, J.W., Opgenoorth, H.J., 2017. On the usage of geomagnetic indices for data selection in internal field modelling. Space Science Reviews 206, 61–90. doi:10.1007/s11214-016-0301-0.

Kelley, M.C., Heelis, R.A., 1989. The Earth's Ionosphere: Plasma Physics and Electrodynamics. volume 43 of *International Geophysics Series*. Academic Press, San Diego.

Kloss, Finlay, 2019. Time-dependent low-latitude core flow and geomagnetic field acceleration pulses. Geophysical Journal International 217, 140–168. doi:10.1093/gji/ggy545.

Kloss, C., Finlay, C.C., Laundal, K.M., Olsen, N., 2023. Polar ionospheric currents and high temporal resolution geomagnetic field models. Geophysical Journal International 235, 1736–1760. doi:10.1093/gji/ggad325.

Kloss, C., Finlay, C.C., Olsen, N., Tøffner-Clausen, L., Gillet, N., Grayver, A., 2026. The CHAOS-8 Geomagnetic Field Model. Earth, Planets and Space 78, 21. URL: https://doi.org/10.1186/s40623-025-02352-0, doi:10.1186/s40623-025-02352-0.

Koelemeijer, P., 2021. Toward Consistent Seismological Models of the Core–Mantle Boundary Landscape. American Geophysical Union (AGU). chapter 9. pp. 229–255. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/9781119528609.ch9, doi:https://doi.org/10.1002/9781119528609.ch9, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/9781119528609.ch9.

Kono, M., 2015. Geomagnetism: An introduction and overview, in: Treatise on Geophysics. 2 ed.. Elsevier.

Kotsiaros, S., et al., 2015. Magnetic field gradient data from swarm satellites. Geophysical Journal International .

Kreutz, B.M., 1984. Mediterranean contributions to the medieval mariner's compass. Technology and Culture 25, 357–394. doi:10.2307/3102323.

Labrosse, S., Poirier, J.P., Le Mouël, J.L., 1997. On cooling of the Earth's core. Physics of the Earth and Planetary Interiors 99, 1–17. URL: https://doi.org/10.1016/S0031-9201(96)03207-4, doi:10.1016/S0031-9201(96)03207-4.

Laj, C., Channell, J.E.T., 2015. Geomagnetic excursions, in: Treatise on Geophysics. 2 ed..

Elsevier.

Langel, R., Ousley, G., Berbert, J., Murphy, J., Settle, M., 1982. The magsat mission. Geophysical Research Letters 9, 243–245. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/GL009i004p00243, doi:https://doi.org/10.1029/GL009i004p00243, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/GL009i004p00243.

Lay, T., Young, C.J., 1990. The stably-stratified outermost core revisited. Geophysical Research Letters 17, 2001–2004. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/GL017i011p02001, doi:https://doi.org/10.1029/GL017i011p02001, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/GL017i011p02001.

Le Mouël, J.L., 1984. Outer-core geostrophic flow and secular variation of earth's geomagnetic field. Nature 311, 734–735. URL: https://doi.org/10.1038/311734a0, doi:10.1038/311734a0.

Lesur, V., Hamoudi, M., Choi, Y., Dyment, J., Thébault, E., 2016. Building the second version of the world digital magnetic anomaly map (wdmam). Earth, Planets and Space 68, 27. doi:10.1186/s40623-016-0404-6.

Lesur, V., Whaler, K.A., Wardinski, I., 2015. Are geomagnetic data consistent with stably stratified flow at the core–mantle boundary? Geophysical Journal International 201, 929–946. URL: https://doi.org/10.1093/gji/ggv031, doi:10.1093/gji/ggv031.

Li, J., Lin, Y., Zhang, K., 2024. Dynamic mode decomposition of the core surface flow inverted from geomagnetic field models. Geophysical Research Letters 51, e2023GL106362. URL: https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2023GL106362, doi:10.1029/2023GL106362.

Li, Y., Zhang, Y., Zhang, T., Yuan, S., Su, W., Sun, J., Feng, F., Cheng, B., Jiang, Y., Liu, P., Xiao, G., Fan, Q., Yu, C., Wang, M., Xiao, Q., Miao, C., 2025. Technical characteristics of macau science satellite-1a. Earth and Planetary Physics 9, 491–499. URL: https://www.eppcgs.org/en/article/doi/10.26464/epp2025062, doi:10.26464/epp2025062.

Livermore, P.W., Hollerbach, R., Finlay, C.C., 2017. An accelerating high-latitude jet in earth's core. Nature Geoscience 10, 62–68. doi:10.1038/ngeo2859.

Lowes, F.J., 1966. Mean-square values on sphere of spherical harmonic vector fields. Journal of Geophysical Research 71, 2179. doi:10.1029/JZ071i008p02179.

Ma, X., Tkalčić, H., 2023. Ccmoc: A new view of the earth's outer core through the global coda correlation wavefield. Physics of the Earth and Planetary Interiors 334, 106957. URL: https://www.sciencedirect.com/science/article/pii/S0031920122001182, doi:https://doi.org/10.1016/j.pepi.2022.106957.

Madsen, F.D., Whaler, K.A., Beggan, C.D., Brown, W.J., Lauridsen, J.B., Holme, R., 2025. Modelling geomagnetic jerks with core surface flow derived from satellite gradient tensor elements of secular variation. Physics of the Earth and Planetary Interiors 366, 107336. URL: https://www.sciencedirect.com/science/article/pii/S0031920125000305, doi:https://doi.org/10.1016/j.pepi.2025.107336.

Mandea, Olsen, 2006. A new approach to directly determine the secular variation from magnetic satellite observations. Geophysical Research Letters 33. URL: https://onlinelibrary.wiley.com/doi/abs/10.1029/2006GL026616, doi:10.1029/2006GL026616.

Mason, S.J., Davies, C.J., Clarke, A.T., Constable, C.G., 2024. Insights into the last 100 ky of geomagnetic field variability using numerical dynamo simulations. Earth and Planetary Science Letters 646, 119011. URL: https://www.sciencedirect.com/science/article/pii/S0012821X24004436, doi:https://doi.org/10.1016/j.epsl.2024.119011.

Masters, T.G., Johnson, S., Laske, G., Bolton, H., 1996. A shear - velocity model of the mantle. Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical

and Engineering Sciences 354, 1385–1411. doi:10.1098/rsta.1996.0054.

Mauersberger, P., 1956. Das Mittel der Energiedichte des geomagnetischen Hauptfeldes an der Erdoberfläche und seine säkulare Änderung. Gerlands Beiträge zur Geophysik 65, 207–215.

Maus, S., 2008. The geomagnetic power spectrum. Geophysical Journal International 174, 135–142. URL: https://doi.org/10.1111/j.1365-246X.2008.03820.x, doi:10.1111/j.1365-246X.2008.03820.x, arXiv:https://academic.oup.com/gji/article-pdf/174/1/135/6076724/174-1-135.pdf.

Maus, S., Weidelt, P., 2004. Separating the magnetospheric disturbance magnetic field into external and transient internal contributions using a 1d conductivity model of the earth. Geophysical Research Letters 31, L12614. doi:10.1029/2004GL020232.

McNamara, A.K., 2019. A review of large low shear velocity provinces and ultra low velocity zones. Tectonophysics 760, 199–220. URL: https://doi.org/10.1016/j.tecto.2018.04.015, doi:10.1016/j.tecto.2018.04.015.

Meier, M.M., Berger, T., Jahn, T., Matthiä, D., Plettenberg, M.C., Scheibinger, M., Schennetten, K., Wirtz, M., 2023. Impact of the south atlantic anomaly on radiation exposure at flight altitudes during solar minimum. Scientific Reports 13, 9348. URL: https://doi.org/10.1038/s41598-023-36190-5, doi:10.1038/s41598-023-36190-5.

Merrill, R.T., McElhinny, M.W., McFadden, P.L., 1996. The Magnetic Field of the Earth: Paleomagnetism, the Core, and the Deep Mantle. volume 63 of *International Geophysics Series*. Academic Press, San Diego.

Metman, Livermore, Mound, Beggan, 2019. Modelling decadal secular variation with only magnetic diffusion. Geophysical Journal International 219, S58–S82. doi:10.1093/gji/ggz089.

Metman, M.C., Beggan, C.D., Livermore, P.W., Mound, J.E., 2020. Forecasting yearly geomagnetic variation through sequential estimation of core flow and magnetic diffusion. Earth, Planets and Space 72, 149. URL: https://doi.org/10.1186/s40623-020-01193-3, doi:10.1186/s40623-020-01193-3.

Minami, T., Nakano, S., Lesur, V., Takahashi, F., Matsushima, M., Shimizu, H., Nakashima, R., Taniguchi, H., Toh, H., 2020. A candidate secular variation model for igrf-13 based on mhd dynamo simulation and 4denvar data assimilation. Earth, Planets and Space 72, 136. URL: https://doi.org/10.1186/s40623-020-01253-8, doi:10.1186/s40623-020-01253-8.

Mosca, I., Cobden, L., Deuss, A., Ritsema, J., Trampert, J., 2012. Seismic and mineralogical structures of the lower mantle from probabilistic tomography. Journal of Geophysical Research: Solid Earth 117. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011JB008851, doi:https://doi.org/10.1029/2011JB008851, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2011JB008851.

Mound, J., Davies, C., Rost, S., Aurnou, J., 2019. Regional stratification at the top of Earth's core due to core–mantle boundary heat flux variations. Nature Geoscience 12, 575–580. URL: https://doi.org/10.1038/s41561-019-0381-z, doi:10.1038/s41561-019-0381-z.

Mound, J.E., Davies, 2020. Scaling laws for regional stratification at the top of earth's core. Geophysical Research Letters 47, e2020GL087715. doi:10.1029/2020GL087715.

Mound, J.E., Davies, 2023. Longitudinal structure of earth's magnetic field controlled by lower mantle heat flow. Nature Geoscience 16, 380–385. doi:10.1038/s41561-023-01148-9.

Munch, F.D., Kemper, J., Khan, A., Helffrich, G., Giardini, D., 2026. Self-consistent models of earth's mantle and core from long-period seismic and tidal constraints. Journal of Geophysical Research: Solid Earth 131, e2024JB030971. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2024JB030971, doi:https://doi.org/10.1029/2024JB030971, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2024JB030971.

e2024JB030971 2024JB030971.

Nakagawa, T., Davies, C.J., 2022. Combined dynamical and morphological characterisation of geodynamo simulations. Earth and Planetary Science Letters 594, 117752. URL: https://doi.org/10.1016/j.epsl.2022.117752, doi:10.1016/j.epsl.2022.117752.

Naskar, S., Mound, J.E., Davies, C.J., Clarke, A.T., 2026. Thermochemical models of outer core convection with heterogeneous core-mantle boundary heat flux. Journal of Studies of Earth's Deep Interior Volume 2. URL: https://jsedi.episciences.org/17084, doi:10.46298/jsedi.17084.

Neubert, T., Mandea, M., Hulot, G., von Frese, R., Primdahl, F., Jørgensen, J.L., Friis-Christensen, E., Stauning, P., Olsen, N., Risbo, T., 2001. Ørsted satellite captures high-precision geomagnetic field data. Eos, Transactions American Geophysical Union 82, 81–88. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/01EO00043, doi:https://doi.org/10.1029/01EO00043, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/01EO00043.

Nilsson, A., Suttie, N., Korte, M., Holme, R., Hill, M., 2020. Persistent westward drift of the geomagnetic field at the core–mantle boundary linked to recurrent high-latitude weak/reverse flux patches. Geophysical Journal International 222, 1423–1432. doi:10.1093/gji/ggaa249.

Ogg, J., 2020. Chapter 5 - geomagnetic polarity time scale, in: Gradstein, F.M., Ogg, J.G., Schmitz, M.D., Ogg, G.M. (Eds.), Geologic Time Scale 2020. Elsevier, pp. 159–192. URL: https://www.sciencedirect.com/science/article/pii/B978012824360200005X, doi:https://doi.org/10.1016/B978-0-12-824360-2.00005-X.

Olsen, N., Friis-Christensen, E., Floberghagen, R., Alken, P., Beggan, C.D., Chulliat, A., Doornbos, E., da Encarnação, J.T., Hamilton, B., Hulot, G., van den IJssel, J., Kuvshinov, A., Lesur, V., Lühr, H., Macmillan, S., Maus, S., Noja, M., Olsen, P.E.H., Park, J., Plank, G., Püthe, C., Rauberg, J., Ritter, P., Rother, M., Sabaka, T.J., Schachtschneider, R., Sirol, O., Stolle, C., Thébault, E., Thomson, A.W.P., Tøffner-Clausen, L., Velímský, J., Vigneron, P., Visser, P.N., 2013. The swarm satellite constellation application and research facility (scarf) and swarm data products. Earth, Planets and Space 65, 1189–1200. URL: https://doi.org/10.5047/eps.2013.07.001, doi:10.5047/eps.2013.07.001.

Olsen, N., Glassmeier, K.H., Jia, X., 2010a. Separation of the magnetic field into external and internal parts. Space Science Reviews 152, 135–157. doi:10.1007/s11214-009-9563-0.

Olsen, N., Hulot, G., Sabaka, T., 2010b. Sources of the Geomagnetic Field and the Modern Data That EnableTheir Investigation. Springer, Berlin, Heidelberg. p. 105–124. URL: https://doi.org/10.1007/978-3-642-01546-5_5, doi:10.1007/978-3-642-01546-5_5.

Olsen, N., Lühr, H., Sabaka, T.J., Mandea, M., Rother, M., Tøffner-Clausen, L., Choi, S., 2006. Chaos—a model of the earth's magnetic field derived from champ, Ørsted, and sac-c magnetic satellite data. Geophysical Journal International 166, 67–75. URL: https://doi.org/10.1111/j.1365-246X.2006.02959.x, doi:10.1111/j.1365-246X.2006.02959.x, arXiv:https://academic.oup.com/gji/article-pdf/166/1/67/5890597/166-1-67.pdf.

Olsen, N., Ravat, D., Finlay, C.C., Kother, L.K., 2017. Lcs-1: a high-resolution global model of the lithospheric magnetic field derived from champ and swarm satellite observations. Geophysical Journal International 211, 1461–1477. doi:10.1093/gji/ggx381.

Olsen, N., Stolle, C., 2017. Magnetic signatures of ionospheric and magnetospheric current systems during geomagnetic quiet conditions—an overview. Space Science Reviews 206, 5–25. doi:10.1007/s11214-016-0279-7.

Olsen, N., et al., 2015. Swarm satellite data and geomagnetic field modelling. Earth, Planets and Space .

Otzen, M., Finlay, C.C., Kloss, C., 2024. Co-estimation of core and lithospheric magnetic

fields by a maximum entropy method. Geophysical Journal International 236, 1646–1667. doi:10.1093/gji/ggae008.

Pais, M.A., Jault, D., 2008. Quasi-geostrophic flows responsible for the secular variation of the earth's magnetic field. Geophysical Journal International 173, 421–443. URL: https://doi.org/10.1111/j.1365-246X.2008.03741.x, doi:10.1111/j.1365-246X.2008.03741.x, arXiv:https://academic.oup.com/gji/article-pdf/173/2/421/6078007/173-2-421.pdf.

Panovska, S., Korte, M., Constable, C.G., 2019. One hundred thousand years of geomagnetic field evolution. Reviews of Geophysics 57, 1289–1337. doi:10.1029/2019RG000656.

Pavón-Carrasco, F.J., Marsal, S., Torta, J.M., Catalán, M., Martín-Hernández, F., Tordesillas, J.M., 2020. Bootstrapping swarm and observatory data to generate candidates for the dgrf and igrf-13. Earth, Planets and Space 72, 152. URL: https://doi.org/10.1186/s40623-020-01198-y, doi:10.1186/s40623-020-01198-y.

Petrov, V.G., Bondar, T.N., 2021. Izmiran candidate field model for igrf-13. Earth, Planets and Space 73, 46. URL: https://doi.org/10.1186/s40623-020-01312-0, doi:10.1186/s40623-020-01312-0.

Pollinger, A., Amtmann, C., Betzler, A., Cheng, B., Ellmeier, M., Hagen, C., Jernej, I., Lammegger, R., Zhou, B., Magnes, W., 2020. In-orbit results of the coupled dark state magnetometer aboard the china seismo-electromagnetic satellite. Geoscientific Instrumentation, Methods and Data Systems 9, 275–291. doi:10.5194/gi-9-275-2020.

Rasson, J.L., 2007. Observatories, Instrumentation. Springer Netherlands, Dordrecht. pp. 711–713. URL: https://doi.org/10.1007/978-1-4020-4423-6_225, doi:10.1007/978-1-4020-4423-6_225.

Rawlinson, N., Pozgay, S., Fishwick, S., 2010. Seismic tomography: A window into deep earth. Physics of the Earth and Planetary Interiors 178, 101–135. URL: https://www.sciencedirect.com/science/article/pii/S0031920109002106, doi:https://doi.org/10.1016/j.pepi.2009.10.002.

Reigber, C., Lühr, H., Schwintzer, P., 2002. Champ mission status. Advances in Space Research 30, 129–134. URL: https://www.sciencedirect.com/science/article/pii/S0273117702002764, doi:https://doi.org/10.1016/S0273-1177(02)00276-4.

Ritsema, J., Lekić, V., 2020. Heterogeneity of seismic wave velocity in Earth's mantle. Annual Review of Earth and Planetary Sciences 48, 377–401. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-earth-082119-065909, doi:10.1146/annurev-earth-082119-065909.

Roberts, P., 2015. 8.03 - theory of the geodynamo, in: Schubert, G. (Ed.), Treatise on Geophysics (Second Edition). second edition ed.. Elsevier, Oxford, pp. 57–90. URL: https://www.sciencedirect.com/science/article/pii/B9780444538024001445, doi:https://doi.org/10.1016/B978-0-444-53802-4.00144-5.

Roberts, P.H., Scott, S., 1965. On analysis of the secular variation. Journal of geomagnetism and geoelectricity 17, 137–151. doi:10.5636/jgg.17.137.

Rogers, H., Davies, C., Constable, C., Mason, S., Frasson, T., 2026. Using paleomagnetic observations and geodynamo simulations to assess earth's magnetic field morphology and variability. Journal of Studies of the Earth's Deep Interior 2. URL: https://doi.org/10.5281/zenodo.21821109, doi:10.5281/zenodo.21821109.

Rogers, H.F., 2022. Improving inversions of regional outer core surface flow models through the application of spherical slepian function URL: https://era.ed.ac.uk/handle/1842/39295, doi:10.7488/era/2546. accepted: 2022-08-02T10:34:07Z.

Rother, M., Korte, M., Morschhauser, A., Vervelidou, F., Matzka, J., Stolle, C., 2021. The mag.num core field model as a parent for igrf-13, and the recent evolution of the south

atlantic anomaly. Earth, Planets and Space 73, 50. URL: https://doi.org/10.1186/s40623-020-01277-0, doi:10.1186/s40623-020-01277-0.

Russell, S., Irving, J.C.E., Myhill, R., Cottaar, S., 2024. The emerging picture of a complex core–mantle boundary. Nature Communications 15. URL: https://doi.org/10.1038/s41467-024-48939-1, doi:10.1038/s41467-024-48939-1.

Sanchez, S., Wicht, J., Bärenzung, J., 2020. Predictions of the geomagnetic secular variation based on the ensemble sequential assimilation of geomagnetic field models by dynamo simulations. Earth, Planets and Space 72, 157. URL: https://doi.org/10.1186/s40623-020-01279-y, doi:10.1186/s40623-020-01279-y.

Schwaiger, T., Gastine, T., Aubert, J., 2019. Force balance in numerical geodynamo simulations: a systematic study. Geophysical Journal International 219, S101–S114. URL: https://doi.org/10.1093/gji/ggz192, doi:10.1093/gji/ggz192.

Schwaiger, T., Jault, D., Gillet, N., Schaeffer, N., Mandea, M., 2023. Local estimation of quasi-geostrophic flows in earth's core. Geophysical Journal International 234, 494–511. URL: https://doi.org/10.1093/gji/ggad089, doi:10.1093/gji/ggad089, arXiv:https://academic.oup.com/gji/article-pdf/234/1/494/50509990/ggad089.pdf.

Stockmann, R., et al., 2009. Slepian functions and localized geomagnetic field modelling. Geophysical Journal International .

Suttie, N., Nilsson, A., Gillet, N., Dumberry, M., 2025. Large-scale palaeoflow at the top of earth's core. Earth and Planetary Science Letters 652, 119185. URL: https://www.sciencedirect.com/science/article/pii/S0012821X24006174, doi:https://doi.org/10.1016/j.epsl.2024.119185.

Tangborn, A., Kuang, W., Sabaka, T.J., Yi, C., 2021. Geomagnetic secular variation forecast using the nasa gems ensemble kalman filter: A candidate sv model for igrf-13. Earth, Planets and Space 73, 47. URL: https://doi.org/10.1186/s40623-020-01324-w, doi:10.1186/s40623-020-01324-w.

Tauxe, L., 2010. Essentials of Paleomagnetism. University of California Press, Berkeley.

Terra-Nova, F., Amit, H., Hartmann, G.A., Trindade, R.I.F., Pinheiro, K.J., 2017. Relating the south atlantic anomaly and geomagnetic flux patches. Physics of the Earth and Planetary Interiors 266, 39–53. doi:10.1016/j.pepi.2017.03.002.

Thébault, E., Purucker, M., Whaler, K.A., Langlais, B., Sabaka, T.J., 2010. The magnetic field of the earth's lithosphere. Space Science Reviews 155, 95–127. doi:10.1007/s11214-010-9667-6.

Usui, Y., Tarduno, J.A., Watkeys, M., Hofmann, A., Cottrell, R.D., 2009. Evidence for a 3.45-billion-year-old magnetic remanence: Hints of an ancient geodynamo from conglomerates of south africa. Geochemistry, Geophysics, Geosystems 10. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2009GC002496, doi:https://doi.org/10.1029/2009GC002496, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2009GC002496.

Valet, J.P., 2003. Time variations in geomagnetic intensity. Reviews of Geophysics 41. doi:10.1029/2001RG000104.

Vine, F.J., 1966. Spreading of the ocean floor: New evidence. Science 154, 1405–1415. doi:10.1126/science.154.3755.1405.

Wardinski, I., Saturnino, D., Amit, H., Chambodut, A., Langlais, B., Mandea, M., Thébault, E., 2020. Geomagnetic core field models and secular variation forecasts for the 13th international geomagnetic reference field (igrf-13). Earth, Planets and Space 72, 155. URL: https://doi.org/10.1186/s40623-020-01254-7, doi:10.1186/s40623-020-01254-7.

Whaler, K., Olsen, N., Finlay, C., 2016. Decadal variability in core surface flows de-

duced from geomagnetic observatory monthly means. Geophysical Journal International 207, 228–243. URL: https://doi.org/10.1093/gji/ggw268, doi:10.1093/gji/ggw268, arXiv:https://academic.oup.com/gji/article-pdf/207/1/228/17370044/ggw268.pdf.

Whaler, K.A., 1980. Does the whole of the earth's core convect? Nature 287, 528–530. doi:10.1038/287528a0.

Whaler, K.A., 1986. Geomagnetic evidence for fluid upwelling at the core-mantle boundary. Geophysical Journal International 86, 563–588. doi:10.1111/j.1365-246X.1986.tb03844.x.

Whaler, K.A., Hammer, M.D., Finlay, C.C., Olsen, N., 2022. Core surface flow changes associated with the 2017 pacific geomagnetic jerk. Geophysical Research Letters 49, e2022GL098616. doi:10.1029/2022GL098616.

Whaler, K.A., Holme, R., 2007. Consistency between the flow at the top of the core and the frozen-flux approximation. Earth, Planets and Space 59, 1219–1229. URL: https://doi.org/10.1186/BF03352070, doi:10.1186/BF03352070.

Wicht, J., Sanchez, S., 2019. Advances in geodynamo modelling. Geophysical & Astrophysical Fluid Dynamics 113, 2–50. URL: https://doi.org/10.1080/03091929.2019.1597074, doi:10.1080/03091929.2019.1597074, arXiv:https://doi.org/10.1080/03091929.2019.1597074.

Yuan, Q., Li, M., 2022. Instability of the african large low-shear-wave-velocity province due to its low intrinsic density. Nature Geoscience 15, 334–339. URL: https://doi.org/10.1038/s41561-022-00908-3, doi:10.1038/s41561-022-00908-3.

Zeren, Z., Yang, Y., Yan, R., Zhang, Z., Wang, J., Huang, H., Xu, S., Lu, H., Zhou, N., Huang, J., 2025. The representative scientific results of the china seismo-electromagnetic satellite. Earthquake Research Advances 5, 100314. URL: https://www.sciencedirect.com/science/article/pii/S277246702400040X, doi:https://doi.org/10.1016/j.eqrea.2024.100314.

Zhao, C., Garnero, E.J., McNamara, A.K., Schmerr, N., Carlson, R.W., 2015. Seismic evidence for a chemically distinct thermochemical reservoir in earth's deep mantle beneath hawaii. Earth and Planetary Science Letters 426, 143–153. URL: https://www.sciencedirect.com/science/article/pii/S0012821X15003684, doi:https://doi.org/10.1016/j.epsl.2015.06.012.

Ziegler, L.B., Constable, C.G., Johnson, C.L., Tauxe, L., 2011. Padm2m: a penalized maximum likelihood model of the 0–2 ma palaeomagnetic axial dipole moment. Geophysical Journal International 184, 1069–1089. URL: https://doi.org/10.1111/j.1365-246X.2010.04905.x, doi:10.1111/j.1365-246X.2010.04905.x, arXiv:https://academic.oup.com/gji/article-pdf/184/3/1069/1731721/184-3-1069.pdf.

# Chapter 2

# Machine Learning

This chapter provides a high-level review of the machine-learning methodology used within this thesis. Section 2.1 introduces the model training and evaluation procedures used throughout Chapters 4–6, including loss functions, optimisation, hyperparameter selection, data scaling and generalisation. Section 2.2 describes the machine-learning architectures used in this thesis, including the neural networks and physics-informed neural networks applied in Chapters 4 and 5, and the convolutional layers and autoencoder classifier used in Chapter 6. Finally, Section 2.3 introduces the explainable artificial intelligence methods used to interpret the classifications in Chapter 6.

A comprehensive introduction to machine learning and deep learning is provided by Goodfellow et al. (2016). More detailed reviews of machine learning in the physical sciences are provided by Mehta et al. (2019); Carleo et al. (2019), physics-informed machine learning by Karniadakis et al. (2021); Cuomo et al. (2022), autoencoders by Bank et al. (2023); Li et al. (2023); Chen and Guo (2023), and explainable artificial intelligence by Linardatos et al. (2021); Kamakshi and Krishnan (2023); Cheng et al. (2025).

Machine learning refers to a set of statistical techniques that leverage data to perform a task, without being explicitly programmed to do so (Armstrong and Fletcher, 2019). This allows relationships and patterns within data to be identified and used to make predictions or draw inferences (Jordan and Mitchell, 2015). Machine-learning methods are particularly useful for large and complex datasets, where the relationship between the available data and the quantity of interest may be difficult to describe analytically. Their use has therefore increased across the physical sciences, where they provide an alternative means of analysing data and investigating problems that may be difficult to address using conventional approaches (Carleo et al., 2019; Mehta et al., 2019).

A machine-learning dataset consists of a collection of examples, where each example represents one observation or item within the dataset (Goodfellow et al., 2016). The information describing an example is supplied to the model as its input, $\boldsymbol{x}$. Depending on the problem, the input may consist of a set of numerical values or structured data, such as an image. A machine-learning model processes this input to produce a predicted output,

$$\hat{\boldsymbol{y}} = f(\boldsymbol{x}; \boldsymbol{\theta}), \tag{2.1}$$

where $\hat{\boldsymbol{y}}$ is the model prediction, $f$ is the function represented by the model, and $\boldsymbol{\theta}$ represents the parameters that determine the mapping between the input and output. These parameters are estimated during training, after which the model can be applied to previously unseen inputs.

Machine learning can be broadly divided into supervised and unsupervised learning. In supervised learning, each input, $\boldsymbol{x}$, is paired with a known target, $\boldsymbol{y}$, which is the numerical value or class that the model is trained to predict. Supervised learning may therefore be used for regression problems, such as predicting house prices from property information, or classification problems, such as determining whether an email is spam. The model parameters are adjusted so that the prediction, $\hat{\boldsymbol{y}}$, approaches the corresponding target across the training examples. In unsupervised learning, no targets are provided and the model instead identifies structure within the inputs. This includes grouping customers with similar purchasing behaviour, identifying unusual financial transactions or obtaining a lower-dimensional representation of a dataset containing many variables (Goodfellow et al., 2016; Carleo et al., 2019; Mehta et al., 2019).

## 2.1 Model Training and Evaluation

A supervised machine-learning model is trained by adjusting its parameters to reduce the difference between its predicted outputs and the corresponding target values. For a training dataset containing $N$ examples, this difference is measured using a loss function, $L$. Training can therefore be described as finding the model parameters that minimise the average loss across the training dataset,

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \frac{1}{N} \sum_{i=1}^{N} L\left(y_i, f(x_i; \boldsymbol{\theta})\right), \tag{2.2}$$

where $\boldsymbol{\theta}^*$ denotes the parameters of the trained model. An optimisation algorithm uses the calculated loss to determine how the model parameters should be updated (Mehta et al., 2019). The loss function defines how the difference between the prediction and target is quantified. Its choice depends on the intended task, since different loss functions may give greater weight to different types or magnitudes of error.

Minimising the loss involves repeatedly applying the model to a set of training examples, comparing its predictions with the corresponding targets, and then updating its parameters. For each batch of examples, the model first produces a set of predictions and the loss function measures their difference from the targets. The optimisation algorithm then determines how the model parameters should be adjusted to reduce this loss. Evaluating one batch and updating the model parameters constitute a training step. The training data may be processed as a single batch or divided into smaller subsets, known as mini-batches, with the parameters updated after each batch. Processing every example in the training dataset once is generally referred to as an *epoch*. However, to avoid confusion with the use of *epoch* to describe a particular time in geomagnetism, this will be referred to as an *iteration* throughout the remainder of this thesis. Training generally continues for multiple iterations until a specified stopping criterion is reached.

### 2.1.1 Loss Functions

Two common supervised-learning tasks are regression and classification. Regression is used to predict a continuous numerical quantity, whereas classification is used to assign an input to one of a set of discrete categories, or to estimate the probability that it belongs to each category (Goodfellow et al., 2016). For regression problems with scalar targets, one commonly used loss

function is the mean squared error (MSE), given by

$$L_{\text{MSE}} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2\,, \tag{2.3}$$

where $y_i$ is the target value and $\hat{y}_i$ is the corresponding model prediction. A benefit of the MSE is that squaring the residuals gives greater weight to large errors, encouraging the model to reduce substantial differences between the predicted and target values. However, this also makes the loss sensitive to outliers, which may have a disproportionate influence on the fitted model.

The Huber loss provides an alternative that is less sensitive to large residuals. It behaves quadratically when the absolute residual is smaller than a specified threshold, $\delta$, and linearly when the residual exceeds this value. For a dataset containing $N$ examples, the mean Huber loss is

$$L_{\text{Huber}} = \frac{1}{N}\sum_{i=1}^{N}\begin{cases}\frac{1}{2}(y_i - \hat{y}_i)^2\,, & |y_i - \hat{y}_i| \leq \delta,\\ \delta\left(|y_i - \hat{y}_i| - \frac{1}{2}\delta\right), & |y_i - \hat{y}_i| > \delta.\end{cases} \tag{2.4}$$

A benefit of the Huber loss is that its quadratic treatment of small residuals encourages accurate predictions, while its linear treatment of residuals larger than $\delta$ reduces the influence of outliers. However, its behaviour depends on the choice of $\delta$, which determines the residual magnitude at which the loss changes from quadratic to linear (Huber, 1964).

For classification problems, the model output is commonly transformed to represent the predicted probability that an input belongs to each possible class. The difference between these probabilities and the known class labels can be measured using the cross-entropy loss. For a classification problem containing $C$ classes, the categorical cross-entropy is given by

$$L_{\text{CE}} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c}\log\left(\hat{y}_{i,c}\right), \tag{2.5}$$

where $y_{i,c}$ indicates whether example $i$ belongs to class $c$, and $\hat{y}_{i,c}$ is the probability assigned to that class by the model. A benefit of cross-entropy is that it rewards the model for assigning a high probability to the correct class and strongly penalises incorrect predictions made with high confidence. However, incorrectly labelled or ambiguous training examples may produce large losses and consequently have a disproportionate influence during training (Mao et al., 2023).

### 2.1.2 Optimisation

An optimiser determines how the model parameters are updated during training in order to reduce the loss function. A widely used optimisation approach is gradient descent. The model first produces predictions from the input data, and the loss function measures the difference between these predictions and the intended training objective. Backpropagation then applies the chain rule to calculate the gradient of the loss with respect to each model parameter, indicating how the loss would change if that parameter were adjusted. The parameters are updated in the direction opposite to the gradient, where the size of the update is controlled by the learning rate. Repeating this process gradually moves the model towards parameter values that produce a lower loss. The parameter update can be written as

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \gamma \nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta}_k), \tag{2.6}$$

where $\boldsymbol{\theta}_k$ contains the parameter values at training step $k$, $\gamma$ is the learning rate, and $\nabla_{\boldsymbol{\theta}} L$ is the gradient of the loss with respect to the model parameters.

For complex models, the loss function may contain several minima (Figure 2.1). A local minimum has a lower loss than neighbouring parameter values, whereas a global minimum has the lowest loss across all possible parameter values. The optimisation may also encounter saddle points and relatively flat regions in which the gradient is small, slowing or preventing further improvements during training (Goodfellow et al., 2016). Gradient-based optimisation is therefore not guaranteed to find the global minimum. However, reaching the global minimum is not always necessary, since different parameter values may produce similarly low losses and comparable model performance.

The solution reached during training may depend on the initial parameter values and other stochastic elements of the training process. Here, stochastic refers to operations that involve random variation and may therefore produce different outcomes each time the model is trained. Many of these operations can be controlled using a random seed, which is a number used to initialise the pseudorandom number generators responsible for parameter initialisation, the ordering of training examples and the composition of training batches. Training the model using several seeds provides different initial conditions and may lead to different solutions, allowing the sensitivity of the model to these choices to be assessed. The best-performing run may then

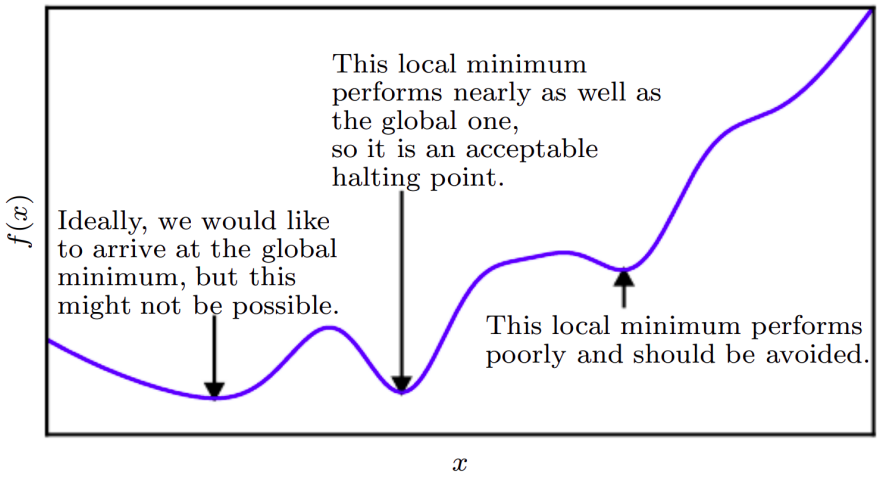



Figure 2.1: Schematic loss function showing a local minimum, which has a lower loss, denoted here by $f(x)$ than the surrounding parameter values, and the global minimum, which has the lowest loss across the complete parameter space. Gradient-based optimisation may converge to the local minimum without reaching the global minimum. Taken from Goodfellow et al. (2016).

be selected, or predictions may be averaged across runs, with their spread used to quantify the variability associated with the training process.

Different optimisers use the gradient in different ways when updating the model parameters. Adam, or adaptive moment estimation, calculates an adaptive update for each parameter using information from the current and previous gradients (Kingma and Ba, 2017). This can improve optimisation where the parameters have gradients of different magnitudes. AdamW extends Adam by applying weight decay separately from the gradient-based update, allowing the strength of this regularisation to be controlled independently (Loshchilov and Hutter, 2017). Both methods are computationally efficient, although neither guarantees that the global minimum will be found. Other extensions of Adam, including SOAP (ShampoO with Adam in the Preconditioner's eigenbasis), have been proposed but are not yet available in stable Python libraries (Vyas et al., 2025). For all the results shown in this thesis, the AdamW optimiser is used.

### 2.1.3 Hyperparameters

Hyperparameters are choices that define the structure of a machine-learning model or control how it is trained, but are not estimated directly by minimising the loss function (Goodfellow et al., 2016). Examples include the learning rate, batch size, number of training iterations and regularisation strength. The learning rate controls the size of each parameter update. A

learning rate that is too small produces only gradual changes and may require many training iterations, whereas one that is too large may repeatedly overshoot regions of lower loss and prevent the optimisation from converging, illustrated schematically in Figure 2.2.

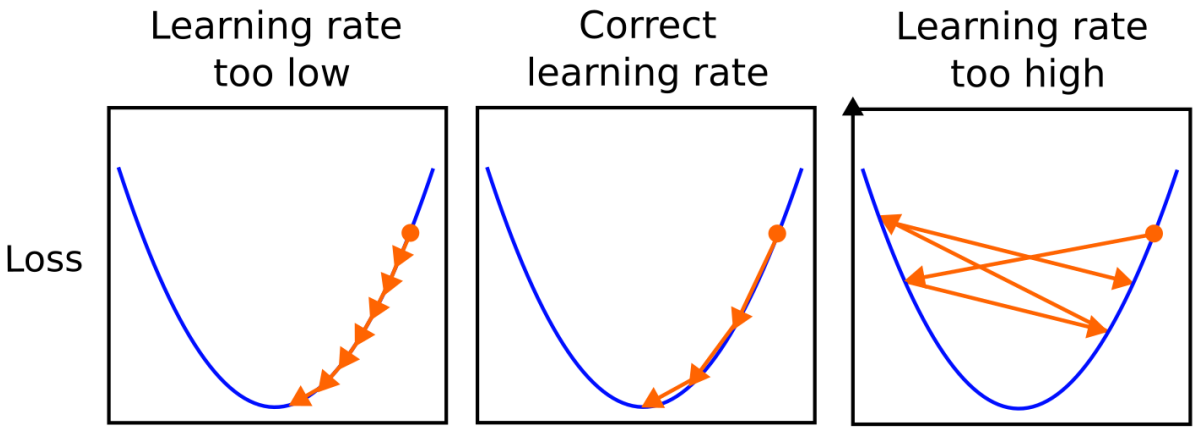



Figure 2.2: Schematic showing the effect of the learning rate on loss minimisation. A small learning rate results in slow convergence, while an excessively large learning rate may prevent convergence.

Hyperparameters therefore affect the rate and stability of training, the complexity of the resulting model and its performance on previously unseen data. Their values may be selected manually by comparing a series of trial models, or automatically using a systematic search method (Bergstra et al., 2012). Manual selection requires an understanding of how each hyperparameter affects model training, whereas automatic approaches, such as Bayesian optimisation, explore the hyperparameter space according to a specified search procedure but generally require additional computation (Snoek et al., 2012). In Chapters 4 and 5, hyperparameter tuning was used to select the neural-network architecture and training parameters for the core-flow inversions, while Chapter 6 investigated the loss weighting, input scaling, maximum spherical harmonic degree and compression ratio of the autoencoder classifier.

The available data are commonly divided into training, validation and test datasets, which are used to estimate the model parameters, select the hyperparameters and provide a final assessment of performance, respectively (Goodfellow et al., 2016). Where several examples originate from the same simulation, experiment or time series, related examples should be kept within the same subset to prevent closely related data from appearing during both training and evaluation. Copies of the model parameters may be saved during training as checkpoints, allowing the model with the lowest validation loss to be retained rather than simply using the final model. This is useful because the optimiser may move away from a previously identified low-

loss solution. As checkpoint and hyperparameter selection use the validation data, a separate test dataset should be retained for the final evaluation where sufficient data are available.

### 2.1.4 Data Scaling

The input and target variables used to train a machine-learning model may have substantially different numerical ranges. These quantities are therefore commonly rescaled before training so that their magnitudes are comparable. For neural networks, appropriate scaling can reduce training time and improve the predictive performance of the resulting model, although performance may depend on the scaling method used (Sola and Sevilla, 1997; Ahsan et al., 2021). Common approaches include normalisation to a specified interval and standardisation by subtracting the mean and dividing by the standard deviation.

The choice of scaling can also affect the information available to the model. Scaling each example independently removes differences in magnitude between examples, whereas applying a common scaling across the dataset retains these differences. The scaling should therefore reflect whether differences in amplitude form part of the relationship that the model is intended to learn. Any scaling parameters estimated from the data, such as the minimum, maximum, mean or standard deviation, should be calculated using the training dataset and then applied unchanged to the validation and test datasets. This maintains the separation between model development and evaluation and prevents information from the evaluation data from influencing the fitted model (Kaufman et al., 2012).

### 2.1.5 Model Evaluation

Once hyperparameter selection and training are complete, the performance of the model can be assessed. The performance measures used depend on the task and are not necessarily the same as the loss function used during training. For regression problems, such as in Chapters 4 and 5, differences between the predicted and target values may be evaluated using measures such as the mean squared error, root mean squared error or mean absolute error (Hodson, 2022). Reconstruction models, which learn to reproduce the original input from a compressed or latent representation, may be evaluated by treating the original input as the target and comparing it with the reconstructed output. For classification problems, such as that considered in Chapter 6, performance may be measured using classification accuracy, which is the proportion of examples assigned to the correct class (Farhadpour et al., 2024). However, accuracy provides

only an overall measure and does not show how well the model performs for each class. This is particularly important for an imbalanced dataset, where high accuracy may be achieved despite poor performance for a class containing relatively few examples. A confusion matrix provides a more detailed assessment by showing the number of examples assigned to each predicted and target class. For the binary example in Figure 2.3, true positives and true negatives are correct predictions, while false positives and false negatives are incorrect predictions. For problems with more than two classes, correct predictions appear along the diagonal, while off-diagonal values show which classes the model has confused (Farhadpour et al., 2024). After model training and selection are complete, these measures should be calculated using a separate test dataset so that the final evaluation is not influenced by data used during model development (Varoquaux and Colliot, 2023).

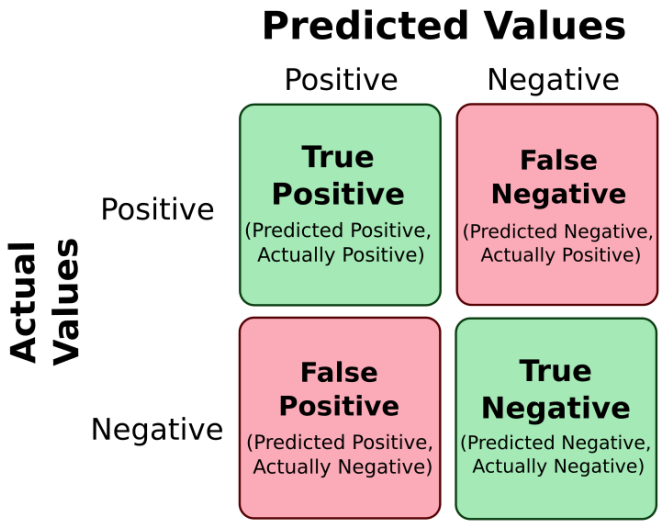



Figure 2.3: Example confusion matrix for a binary classification problem. Correct predictions are shown by the true-positive and true-negative values, while false positives and false negatives indicate the two types of incorrect classification.

### 2.1.6 Generalisation

The purpose of training a machine-learning model is to learn a relationship that can be applied to previously unseen data, rather than simply reproducing the training examples. This ability is known as generalisation (Goodfellow et al., 2016; Bejani and Ghatee, 2021). A model underfits when it is unable to represent the relationship within the training data, resulting in poor performance on both the training and validation datasets. In contrast, overfitting occurs when a model performs well on the training data but poorly on the validation data because the fitted relationship does not generalise to new examples (Bejani and Ghatee, 2021). Whether a model underfits or overfits depends partly on its capacity, which describes the complexity of the relationships that it can represent. A model with insufficient capacity may be unable to represent the underlying structure of the data, whereas increasing the capacity can allow it to fit patterns that are specific to the training examples (Belkin et al., 2019). The relationship between capacity and generalisation is not always straightforward, however, as highly overparameterised models can sometimes fit the training data closely while still performing well on unseen examples (Belkin et al., 2019). Examples of underfitting, overfitting and an appropriate model capacity are shown in Figure 2.4.

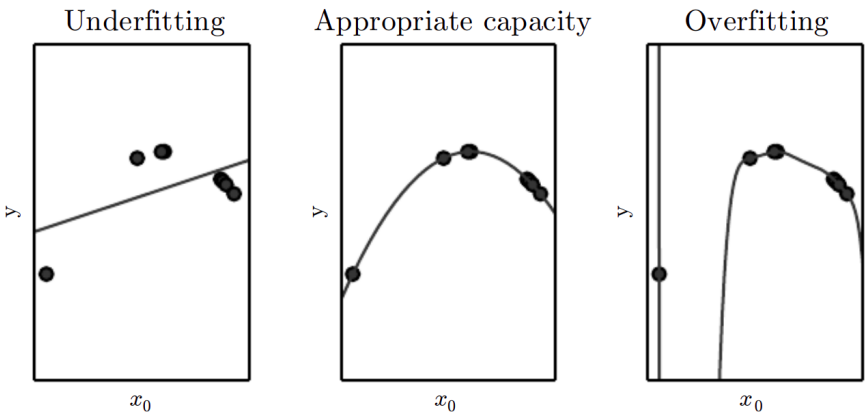



Figure 2.4: Examples of underfitting, an appropriate model capacity and overfitting for a quadratic dataset. The overfitted model reproduces the training examples closely but does not recover the underlying relationship. Taken from Goodfellow et al. (2016).

Generalisation can be improved by selecting the model structure and hyperparameters using validation data, applying regularisation or early stopping, and ensuring that the training data are sufficiently representative of the intended application. Early stopping provides another form of regularisation by ending training when performance on the validation dataset no longer improves. Regularisation refers to modifications to the model or training procedure that restrict

the fitted solution. These include penalties on the magnitude of the model parameters (Goodfellow et al., 2016), dropout, which randomly excludes a proportion of network units during training (Srivastava et al., 2014), and restrictions on the size of the network. This is separate from, but can be combined with, regularisation applied to diagnostic properties of the solution, such as the kinetic energy of the flow.

The similarity between the data used for training and those encountered when the model is applied is important. Models that perform well on data drawn from the training distribution may be sensitive to distribution shifts, where the statistical properties of the application data differ from those represented during training (Krueger et al., 2021; Klivans et al., 2024). Predictions may also require extrapolation when the input lies beyond the range of the training data, where the behaviour of the model is less well constrained by the available examples (Cao and Yousefzadeh, 2023). This is particularly relevant when a model is trained using synthetic data or numerical simulations and subsequently applied to observations, since the simulated system provides only an approximation of the system encountered during application and may introduce a distribution shift (Wagenmaker et al., 2024). Model performance should therefore be evaluated using test data that represent the intended application, while predictions beyond the range or distribution of the training data should be identified and interpreted with care.

## 2.2 Machine-Learning Architectures

The architecture of a machine-learning model describes how its components are arranged and how information passes between them, and has an important influence on the relationships that the model can represent. One challenge is constructing models that can identify complex features directly from raw data. This can be addressed using deep learning, in which multiple layers transform the input through a sequence of simpler non-linear mappings. Early layers identify simple features, while later layers combine these into more complex representations of the data, allowing relationships involving several input features to be learned (LeCun et al., 2015). This section introduces the deep-learning architectures used within this thesis, beginning with feed-forward neural networks before considering physics-informed neural networks, convolutional layers and autoencoder classifiers.

### 2.2.1 Neural Networks

Feed-forward neural networks are a common class of deep-learning model in which information passes from the input to the output without feedback connections (Goodfellow et al., 2016). Neural networks were originally inspired by biological neural systems, with interconnected artificial neurons providing a simplified representation of how the brain processes information. Their ability to represent complex relationships is supported by the universal approximation theorem, which states that a feed-forward neural network with a sufficiently large hidden layer can approximate any continuous function on a bounded domain to arbitrary accuracy (Cybenko, 1989; Hornik et al., 1989). As shown in Figure 2.5, a feed-forward neural network consists of an input layer, one or more hidden layers and an output layer. The input layer receives the information describing each example, while each hidden layer transforms the output of the preceding layer. The output layer then uses the final hidden-layer representation to produce the model prediction. The network therefore represents the overall mapping as a sequence of transformations, with the output of each layer providing the input to the next.

Each hidden layer contains a set of computational units known as neurons. For neuron $j$, the inputs from the preceding layer are multiplied by their corresponding weights and combined with a bias,

$$z_j = \sum_{i=1}^{n} w_{ji} x_i + b_j, \tag{2.7}$$

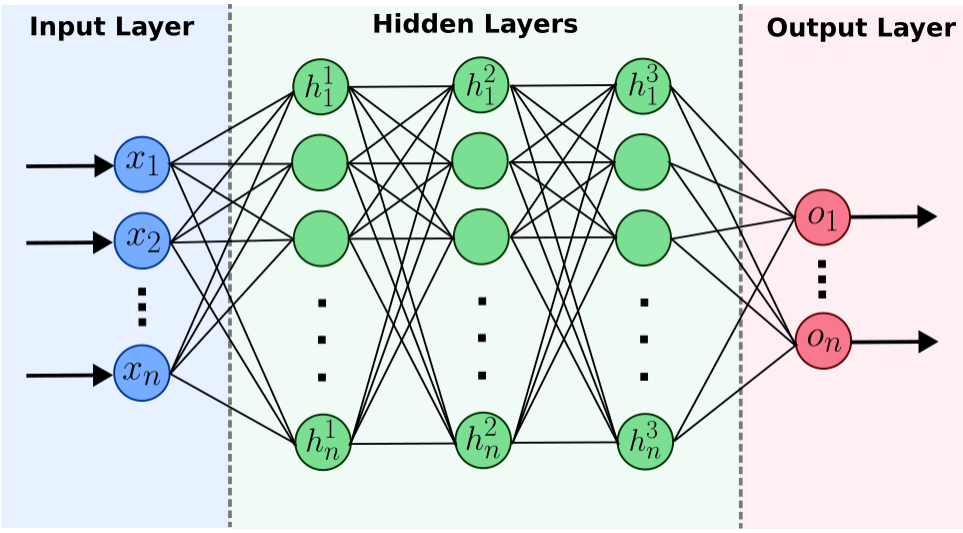



Figure 2.5: Example schematic of a neural network, showing the input layer, hidden layers and output layer.

where $x_i$ is the output of neuron $i$ in the preceding layer, $w_{ji}$ is the weight connecting neuron $i$ to neuron $j$, $b_j$ is the bias associated with neuron $j$, and $z_j$ the weighted input of neuron $j$. An activation function, $a$, is then applied to this weighted input to give the output of the neuron,

$$h_j = a(z_j). \tag{2.8}$$

The weights determine how strongly each input contributes to the neuron, while the bias shifts the value to which the activation function is applied. This allows the neuron to produce a non-zero response when its inputs are zero and changes the input value at which the neuron becomes active, increasing the range of relationships that the network can represent. Before training, the weights are initialised, typically using small random values, to provide a starting point for optimisation. During training, the weights and biases are adjusted iteratively to reduce the loss between the model predictions and corresponding targets.

The same calculation is performed by every neuron within a layer and can therefore be written in matrix form. The output vector from the first hidden layer is

$$\boldsymbol{h}^{(1)} = a^{(1)}\left(\boldsymbol{W}^{(1)}\boldsymbol{x} + \boldsymbol{b}^{(1)}\right), \tag{2.9}$$

where $\boldsymbol{x}$ is the input to the hidden layer, $\boldsymbol{W}^{(1)}$ contains the weights connecting the input to the first hidden layer, $\boldsymbol{b}^{(1)}$ contains the biases of its neurons, and $a^{(1)}$ is the activation function applied to the resulting vector. The superscript identifies the layer and does not represent a

power. The output of the first hidden layer then becomes the input to the second hidden layer,

$$\boldsymbol{h}^{(2)} = a^{(2)}\left(\boldsymbol{W}^{(2)}\boldsymbol{h}^{(1)} + \boldsymbol{b}^{(2)}\right). \tag{2.10}$$

The same operation is repeated through the remaining hidden layers, with the output of each layer providing the input to the next.

The size of a feed-forward neural network is determined by its depth and width. The depth gives the number of successive layers, while the width gives the number of neurons within each hidden layer. Neurons within the same layer learn different combinations of the outputs from the preceding layer, while successive layers combine these outputs into progressively more complex representations. Increasing the width allows a layer to represent more features in parallel, whereas increasing the depth allows those features to be combined through a greater number of transformations. Depth and width can therefore both increase the capacity of the network, but they affect the learned representation in different ways (Raghu et al., 2017; Lu et al., 2017). There is no agreed upon 'best ratio' between them, so the architecture is usually selected by comparing networks with different depths and widths using validation data. A network with insufficient capacity may be unable to represent the required relationship, whereas unnecessary complexity can make training more difficult and may allow the model to learn features that do not generalise to new examples (Bejani and Ghatee, 2021; He et al., 2016). The selected architecture should therefore provide a balance between validation performance, training stability and model complexity.

The activation function determines how the weighted input to each neuron is transformed before being passed to the next layer. Without non-linear activation functions, a sequence of fully connected layers would remain equivalent to a single linear transformation, irrespective of the number of layers. Non-linear activation functions therefore allow neural networks to represent non-linear relationships between their inputs and outputs (Goodfellow et al., 2016). Common activation functions are shown in Table 2.1. The appropriate activation function depends on the architecture, the position of the layer and the intended form of the model output.

| **Activation Function** | **Equation** | **Output Range** | **Notes** |
|---|---|---|---|
| Sigmoid | $a(z) = \dfrac{1}{1+\exp(-z)}$ | $(0, 1)$ | Suitable for binary and multilabel output layers. Gradients become small for large-magnitude inputs, which can slow training. |
| Hyperbolic tangent | $a(z) = \tanh(z)$ | $(-1, 1)$ | Produces outputs centred about zero, but gradients become small for large-magnitude inputs. |
| ReLU | $a(z) = \max(0, z)$ | $[0, \infty)$ | Computationally efficient and does not saturate for positive inputs. Neurons receiving consistently negative inputs may cease to update. |
| Swish | $a(z) = z\,\mathrm{sigmoid}(\beta z)$ | Bounded below and unbounded above | Smooth and retains a non-zero response for some negative inputs, but is more computationally expensive than ReLU. |

Table 2.1: Summary of commonly used neural-network activation functions (Goodfellow et al., 2016; Ramachandran et al., 2017). Here, $z$ is the weighted input to a neuron.

### 2.2.2 Physics-Informed Neural Networks

Physics-informed neural networks (PINNs) are a class of neural network that incorporate mathematical descriptions of physical processes into model training (Raissi et al., 2019). In a conventional neural network, the model parameters are estimated by minimising the difference between the predicted and target values. In a PINN, the loss function also measures the extent to which the prediction satisfies the governing physical equations. This allows observations and prior physical knowledge to be combined within a single model, restricting the range of acceptable solutions to those that are consistent with both the available data and the imposed physical constraints.

Physical information is incorporated by including a physics-based component in the loss function. This component measures how closely the network prediction satisfies the governing equations by calculating the residual produced when the prediction is substituted into them. The total loss may therefore be written as

$$L_{\text{total}} = \alpha_{\text{data}} L_{\text{data}} + \alpha_{\text{physics}} L_{\text{physics}}, \tag{2.11}$$

where $L_{\text{data}}$ measures the difference between the predicted and observed values, and $L_{\text{physics}}$ measures the extent to which the prediction satisfies the governing equations. The weighting parameters $\alpha_{\text{data}}$ and $\alpha_{\text{physics}}$ control the relative importance of these two objectives during

training. They do not need to sum to one, and a term may be omitted when the corresponding constraint is not required.

Where the governing equations contain derivatives of the predicted quantities, these derivatives are calculated using automatic differentiation. A neural network represents its output through a sequence of mathematical operations. Automatic differentiation records this sequence and applies the chain rule to calculate the derivative of the network output with respect to its input coordinates (Raissi et al., 2019; Haghighat and Juanes, 2021). For example, if a network predicts a quantity $u$ as a function of the spatial coordinates $\boldsymbol{x}$, automatic differentiation calculates derivatives such as

$$\nabla_{\boldsymbol{x}} u = \frac{\partial u}{\partial \boldsymbol{x}}. \tag{2.12}$$

The derivative is calculated directly from the operations defining the network, rather than by comparing its output at neighbouring points. This avoids the step-size and truncation errors associated with finite-difference approximations and gives the derivative of the network function up to floating-point rounding error (Baydin et al., 2018). However, the accuracy of this derivative as a representation of the physical solution still depends on how accurately the network represents that solution.

The calculated derivatives are substituted into the governing equations, and the resulting residual forms the physical component of the loss function. This residual may be evaluated at the observation locations or at additional locations distributed throughout the solution domain, known as collocation points. Collocation points allow the governing equations to constrain the solution where observations are not available (Raissi et al., 2019; Cuomo et al., 2022). The governing equations are commonly imposed as soft constraints, meaning that a non-zero residual may remain after training. Some conditions may instead be satisfied exactly through the construction of the network output (Karniadakis et al., 2021; Cuomo et al., 2022). In the present context, the tangentially geostrophic constraint in (1.31) reduces the radial induction equation to the form given in Equation (1.34). Since the PINN predicts only the tangential flow components, this condition is satisfied by construction and does not require an additional loss term.

The inclusion of physical information restricts the range of acceptable solutions to those that are largely consistent with the governing equations. This can reduce ambiguity when observations

are sparse or noisy and may reduce the training data required compared with a purely data-driven model (Raissi et al., 2019; Karniadakis et al., 2021). Evaluating the physical loss at additional collocation points allows the governing equations to constrain the solution where observations are unavailable. As these points do not require the fixed spatial mesh used by many conventional numerical methods, they can be distributed flexibly across irregular domains. Once trained, the network provides a continuous and differentiable approximation that can be evaluated throughout the solution domain.

The performance of a PINN depends partly on the weighting of the different loss terms, since an imbalance may cause the optimisation to prioritise one objective over another (Karniadakis et al., 2021; Cuomo et al., 2022). Training may also become more challenging for problems containing sharp gradients or processes operating over different spatial or temporal scales. The predictions additionally depend on the validity of the physical equations and assumptions imposed during training. These considerations require appropriate network design and validation, but PINNs remain a flexible method for combining observations with physical constraints across a range of scientific problems.

PINNs have been applied to a range of scientific problems. In fluid mechanics, PINNs have been used to infer velocity and pressure fields and to solve problems governed by the Navier–Stokes equations (Raissi et al., 2020; Sharma et al., 2023). Applications in medicine include reconstructing blood-flow patterns from limited observations (Arzani et al., 2021), while applications in nuclear engineering include neutron transport (Schiassi et al., 2022). Livermore et al. (2024) used PINNs to reconstruct Jupiter's internal magnetic field from sparse Juno measurements, obtaining a similar fit to the data and similar magnetic-field structure at and above the surface to existing methods, while allowing weak ambient electrical currents and avoiding the amplification of small-scale noise during inward continuation. Within geophysics, Chen et al. (2022) used the eikonal equation within a neural network for seismic tomography, obtaining results consistent with conventional methods while reducing sensitivity to travel-time errors. More recently, PINNs have been applied to infer core-surface flow from geomagnetic field models. Li et al. (2024) incorporated the magnetic induction equation, a tangentially geostrophic flow constraint and unresolved small-scale magnetic fields within a PINN inversion, before applying dynamic mode decomposition to the recovered flow. Li et al. (2025) extended this approach by incorporating quasi-geostrophic constraints and separate network representations for the flow

inside and outside the tangent cylinder.

### 2.2.3 Convolutional Layers

The analysis in Chapter 6 uses two-dimensional images of the radial magnetic field as input to the autoencoder classifier. Images contain many input values arranged in a two-dimensional spatial structure. Processing every image value using a fully connected layer can require a large number of trainable parameters and become computationally expensive, particularly for high-resolution images (LeCun and Bengio, 1998; Masci et al., 2011). Convolutional layers provide a more efficient approach by processing small local regions and sharing the same weights across different image locations. This preserves relationships between neighbouring values and allows the network to learn spatial features, such as edges and increasingly complex patterns (LeCun et al., 2015; Masci et al., 2011). Convolutional layers are therefore widely used for image classification, reconstruction and feature extraction (Krizhevsky et al., 2012; Masci et al., 2011). Unlike a fully connected layer, in which every input is connected to every output, a convolutional layer processes small local regions of the input using a kernel. A kernel is a small array of trainable weights that moves across the input and calculates a weighted combination of the values within each region (Goodfellow et al., 2016). A channel is one layer of information stored at each position in an image. For the magnetic-field images used in Chapter 6, each grid point contains only the radial magnetic field, $B_r$, so the input has a single channel. In comparison, a colour image generally has three channels containing the red, green and blue values at each pixel. For a single-channel two-dimensional input, $I$, and kernel, $K$, the convolution operation is commonly written as

$$S(i,j) = (I \star K)(i,j) + b = \sum_m \sum_n I(i+m, j+n) K(m,n) + b, \tag{2.13}$$

where $S$ is the resulting feature map and $b$ is the trainable bias associated with the kernel. Here, $i$ and $j$ identify the row and column of the output feature map, while $m$ and $n$ identify positions within the kernel and the corresponding local region of the input. This process is shown in Figure 2.6. Although this operation is technically cross-correlation because the kernel is not reversed, it is commonly referred to as convolution in neural-network libraries (Goodfellow et al., 2016).

A convolutional layer generally contains several kernels, each with a different set of trainable

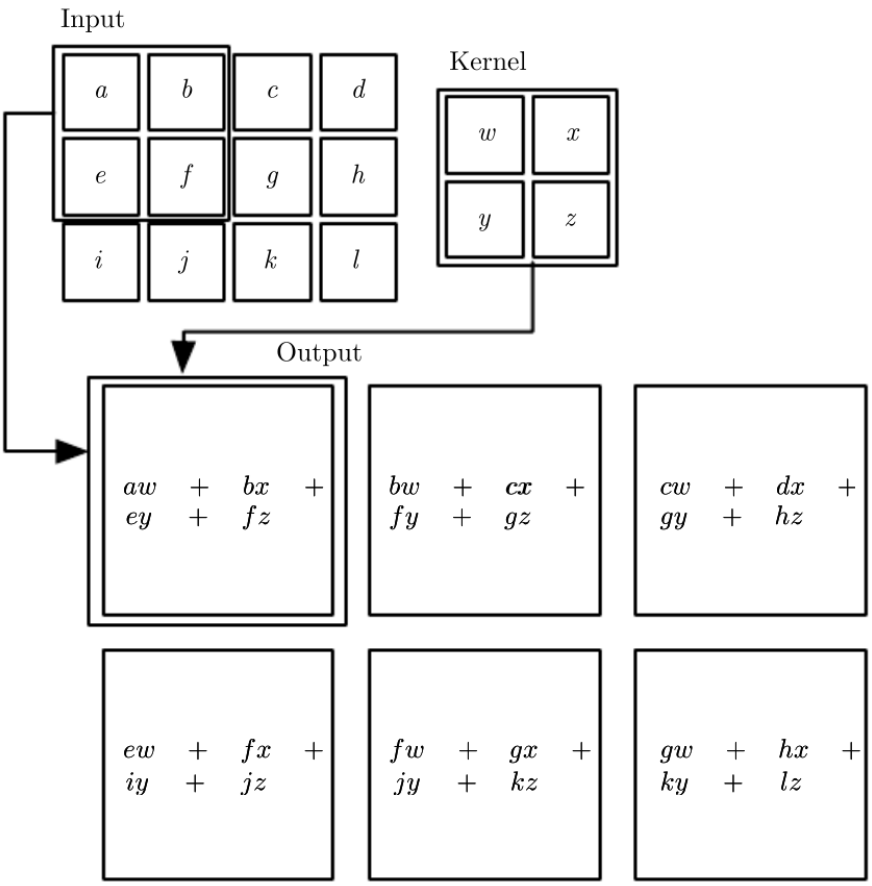



Figure 2.6: Example of two-dimensional cross-correlation without kernel flipping. The output is calculated only where the kernel lies entirely within the input, with no padding applied. Taken from Goodfellow et al. (2016).

weights. Each kernel moves across the input and produces a feature map showing its response at each location. Different kernels can therefore learn to identify different spatial patterns, with the resulting feature maps forming the output channels of the layer. The same kernel weights are applied at every location, allowing a pattern to be detected wherever it occurs and reducing the number of trainable parameters. Successive layers combine these local patterns, allowing the network to identify features that extend across increasingly large regions of the original input (Goodfellow et al., 2016; Araujo et al., 2019).

The dimensions of the resulting feature map are controlled by the kernel size, stride and padding. The kernel size determines the extent of the local region processed at each position, while the stride gives the number of grid points by which the kernel moves between successive applications. A larger stride produces a smaller feature map because the kernel is evaluated at fewer locations. Padding adds values around the boundary of the input, allowing features near its edges to contribute to more kernel positions and controlling the reduction in spatial dimensions. An

activation function is commonly applied after convolution, followed in some architectures by a pooling operation that replaces values within a local region with a summary statistic, such as their maximum or mean. Pooling reduces the spatial dimensions and computational cost of subsequent layers, while making the representation less sensitive to small changes in feature location. However, pooling may remove useful information where the precise spatial position of a feature is important (Goodfellow et al., 2016).

### 2.2.4 Autoencoder Classifiers

An autoencoder is a neural network that learns to reconstruct its input after passing it through a compressed representation (Bank et al., 2023; Li et al., 2023). It consists of two components. The encoder maps the input to a latent representation, which is a compressed set of features learned from the input data rather than specified directly. The decoder then maps this representation back to the original input space to produce a reconstruction. The encoder and decoder are trained together by minimising a reconstruction loss that measures the difference between the input data and the reconstructed data. A schematic for this is shown in Figure 2.7.

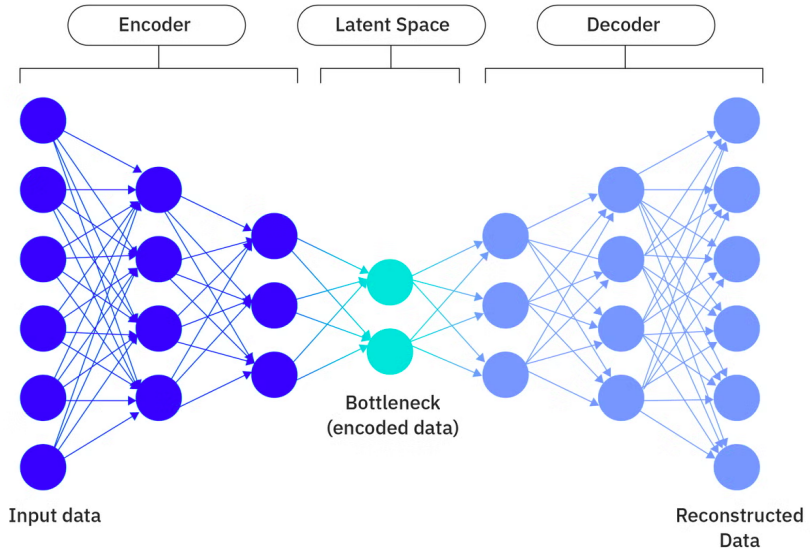



Figure 2.7: Schematic of an autoencoder, taken from Bergmann and Stryker (2026).

The latent representation forms a bottleneck within the network. Here, latent refers to features that are learned from the input data rather than observed directly. If the latent representation contains fewer values than the input, the encoder reduces the dimensionality of the data by representing it using a smaller set of learned features. A smaller latent representation results in greater compression, but may also discard information required for reconstruction or

classification.

An autoencoder must be constrained to prevent it from simply copying the input without learning a useful representation. A lower-dimensional bottleneck provides this constraint by forcing the encoder to preserve only the information most important for reconstructing the input. This information is stored in the latent feature maps, which form the compressed representation. If the bottleneck is too small, important information may be lost, whereas if it is too large, the network may reproduce the input without learning a meaningful compressed representation (Cheng et al., 2020; Chen and Guo, 2023).

A classifier trained without additional architecture may learn features that separate the training classes without retaining the broader structure of the input, including features that are specific to the training examples and may not generalise to new data. Connecting a classification head, the output layer used for class prediction, to an autoencoder provides an additional reconstruction objective, requiring the latent representation to support both classification and reconstruction. The classification head uses the encoded features to predict the class, while the decoder requires them to retain sufficient information to reconstruct the input. Training may therefore combine a reconstruction loss, $L_{\mathrm{r}}$, and a classification loss, $L_{\mathrm{c}}$,

$$L_{\mathrm{total}} = \alpha_{\mathrm{r}} L_{\mathrm{r}} + \alpha_{\mathrm{c}} L_{\mathrm{c}}, \tag{2.14}$$

where $\alpha_{\mathrm{r}}, \alpha_{\mathrm{c}}$ control the relative contributions of the two objectives. The reconstruction loss encourages the latent representation to retain information describing the overall structure of the input, while the classification loss encourages it to retain features that distinguish between the target classes. The model is therefore trained to preserve information that supports both reconstruction and classification, although the relative weighting is important because features that are useful for one task may not be useful for the other.

Autoencoder classifiers have been applied in fields including medical-image classification (Ganotra et al., 2024; Mao et al., 2025), anomaly detection (Saeedi and Giusti, 2023; Su et al., 2015; Kozmin et al., 2025), hyperspectral imaging (Xing et al., 2016) and remote-sensing image classification (Cheng et al., 2020). Variants of the conventional autoencoder have been developed to encourage the extraction of more useful features. Denoising autoencoders learn to reconstruct the original input after it has been perturbed, while sparse autoencoders limit the number of

latent features that can be active at the same time. These constraints reduce the model's ability to simply reproduce the input and encourage it to learn the underlying structure of the data (Li et al., 2023; Chen and Guo, 2023).

## 2.3 Explainable Artificial Intelligence

Explainable artificial intelligence (XAI) encompasses methods that make the outputs of machine-learning models more understandable, helping to reduce the so-called 'black-box' problem (Doshi-Velez and Kim, 2017; Linardatos et al., 2021). These methods can identify the features influencing a prediction, investigate incorrect classifications and help assess whether a model is responding to expected features or to unintended correlations within the training data, such as image artefacts that are associated with only one class (Kamakshi and Krishnan, 2023; Cheng et al., 2025). This is particularly important where the basis of an automated decision must be understood, including in fields such as healthcare, public policy and science.

Although the terms are not used consistently across the literature, interpretability commonly refers to how readily the operation of a model can be understood, while explainability methods aim to clarify the basis of its predictions (Doshi-Velez and Kim, 2017). The internal relationships of a simple model may be examined directly, but this becomes more difficult as model complexity increases. Explainability methods instead provide an approximation of how the input contributed to the model output. SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017) and Local Interpretable Model-Agnostic Explanations (LIME) (Ribeiro et al., 2016) estimate the contribution of individual input variables, while saliency maps provide spatial explanations for image-based models (Linardatos et al., 2021; Kamakshi and Krishnan, 2023; Cheng et al., 2025). The appropriate method depends on the form of the input and whether the aim is to understand the overall model or explain an individual prediction.

A saliency map provides an indication of which parts of an input image matter most when the model makes a particular classification, and can therefore be interpreted as showing where the model is "looking" when making its decision (Linardatos et al., 2021; Cheng et al., 2025). For an input image $\boldsymbol{x}$, the classifier produces one numerical output for each possible class. Let $s_c(\boldsymbol{x})$ denote the output for class $c$, where a larger value indicates greater support for assigning the

image to that class. A basic gradient saliency map calculates

$$S_i = \left| \frac{\partial s_c(\boldsymbol{x})}{\partial x_i} \right|, \tag{2.15}$$

where $x_i$ is the value of pixel $i$ and $S_i$ is its saliency. This derivative measures how much the output for class $c$ would change following a small change in that pixel. Pixels with larger values of $S_i$ are therefore locations where a small change would produce a larger local change in the class output. Calculating $S_i$ across the image produces a map showing the regions to which the classification is most sensitive. The mathematical formulation is provided by Simonyan et al. (2014). However, the map shows where changes would most affect the classification and does not demonstrate that a region physically causes the model to select that class.

Several extensions have been developed to improve the clarity and stability of gradient-based saliency maps. SmoothGrad calculates saliency maps for several slightly perturbed versions of the input and averages the results, reducing visual noise in the resulting sensitivity map (Smilkov et al., 2017). Gradient-weighted class activation mapping (Grad-CAM) provides a broader indication of where the model is "looking" when it assigns an image to a particular class. Rather than measuring the sensitivity of the classification to each input pixel directly, Grad-CAM examines the feature maps produced within a convolutional layer. These feature maps represent spatial patterns identified by the model. Grad-CAM measures how strongly each feature map affects the output for the selected class and combines the maps according to their influence. The resulting map highlights regions that provide positive support for the selected class (Selvaraju et al., 2017; Linardatos et al., 2021; Cheng et al., 2025). An example is shown in Figure 2.8, where the highlighted regions show which parts of each image contributed most strongly to the model's answer. Since the feature maps have a lower spatial resolution than the original image, Grad-CAM identifies broader influential regions rather than individual pixels. This makes the results easier to interpret spatially, but provides coarser spatial localisation than a pixel-level saliency map.

Saliency maps provide an intuitive means of identifying the spatial regions associated with a prediction and assessing whether a model responds to the expected features. However, gradient-based maps measure the local sensitivity of the output to changes in the input rather than demonstrating a causal relationship, and their appearance may depend on the selected output class, network layer, smoothing procedure and attribution method (Linardatos et al., 2021;

Kamakshi and Krishnan, 2023; Cheng et al., 2025). When combined with physical knowledge and appropriate validation, saliency maps provide a useful means of determining which spatial features influence a model and whether its predictions are based on scientifically meaningful structure.

Explainable artificial intelligence has been applied in fields including medical-image analysis, where explanations can show which image regions influenced a diagnosis (Muhammad and Bendechache, 2024), remote sensing, where it has been used to interpret image classifications (Höhl et al., 2024), and autonomous driving, where explanations can support the monitoring and validation of safety-critical decisions (Kuznietsov et al., 2024).

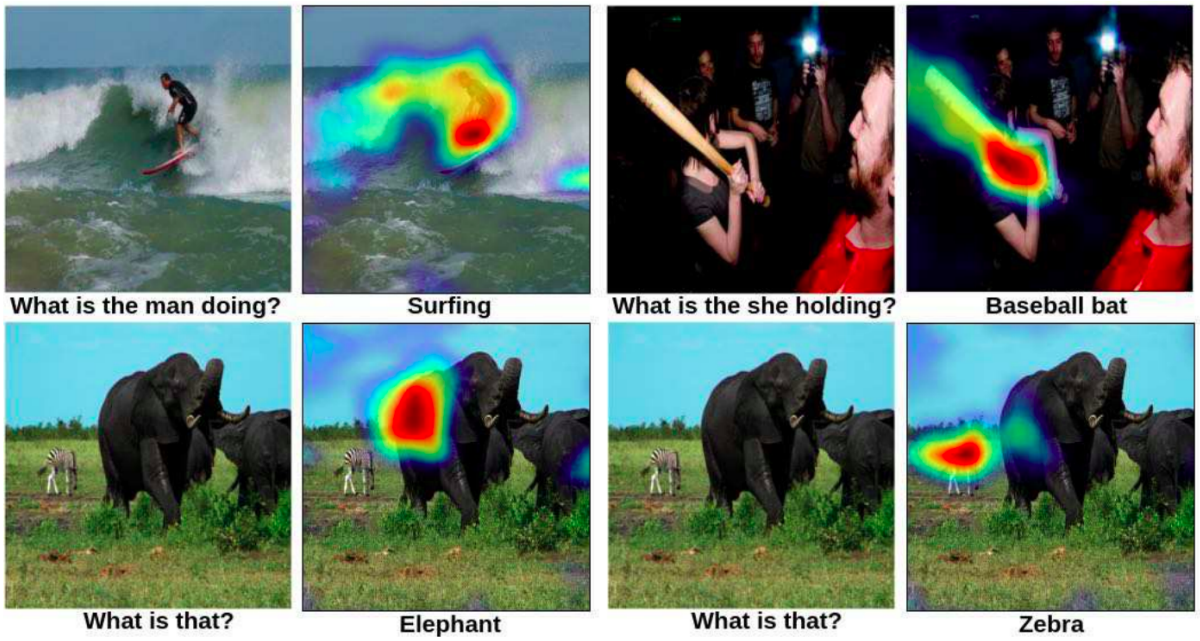



Figure 2.8: Example Grad-CAM visualisations for a visual-question-answering model. The highlighted regions indicate the parts of each image that contributed most strongly to the model's response to the corresponding question. Taken from Selvaraju et al. (2017).

# References


Ahsan, M.M., Mahmud, M.A.P., Saha, P.K., Gupta, K.D., Siddique, Z., 2021. Effect of data scaling methods on machine learning algorithms and model performance. Technologies 9, 52. doi:10.3390/technologies9030052.

Araujo, A., Norris, W., Sim, J., 2019. Computing receptive fields of convolutional neural networks. Distill doi:10.23915/distill.00021.

Armstrong, J.A., Fletcher, L., 2019. Fast solar image classification using deep learning and its importance for automation in solar physics. Solar Physics 294, 80. doi:10.1007/s11207-019-1473-z.

Arzani, A., Wang, J.X., D'Souza, R.M., 2021. Uncovering near-wall blood flow from sparse data with physics-informed neural networks. Physics of Fluids 33, 071905. doi:10.1063/5.0055600.

Bank, D., Koenigstein, N., Giryes, R., 2023. Autoencoders. Springer International Publishing, Cham. pp. 353–374. URL: https://doi.org/10.1007/978-3-031-24628-9_16, doi:10.1007/978-3-031-24628-9_16.

Baydin, A.G., Pearlmutter, B.A., Radul, A.A., Siskind, J.M., 2018. Automatic differentiation in machine learning: A survey. Journal of Machine Learning Research 18, 1–43. URL: https://jmlr.org/papers/v18/17-468.html.

Bejani, M.M., Ghatee, M., 2021. A systematic review on overfitting control in shallow and deep neural networks. Artificial Intelligence Review 54, 6391–6438. doi:10.1007/s10462-021-09975-1.

Belkin, M., Hsu, D., Ma, S., Mandal, S., 2019. Reconciling modern machine-learning practice and the classical bias–variance trade-off. Proceedings of the National Academy of Sciences 116, 15849–15854. doi:10.1073/pnas.1903070116.

Bergmann, D., Stryker, C., 2026. What is a variational autoencoder? IBM Think. URL: https://www.ibm.com/think/topics/variational-autoencoder. accessed 11 August 2026.

Bergstra, J., Yamins, D., Cox, D.D., 2012. Making a science of model search URL: http://arxiv.org/abs/1209.5111, doi:10.48550/arXiv.1209.5111. arXiv:1209.5111 [cs].

Cao, X., Yousefzadeh, R., 2023. Extrapolation and AI transparency: Why machine learning models should reveal when they make decisions beyond their training. Big Data & Society 10. doi:10.1177/20539517231169731.

Carleo, G., Cirac, I., Cranmer, K., Daudet, L., Schuld, M., Tishby, N., Vogt-Maranto, L., Zdeborová, L., 2019. Machine learning and the physical sciences. Reviews of Modern Physics 91, 045002. doi:10.1103/RevModPhys.91.045002.

Chen, S., Guo, W., 2023. Auto-encoders in deep learning: A review with new perspectives. Mathematics 11, 1777. doi:10.3390/math11081777.

Chen, Y., de Ridder, S.A.L., Rost, S., Guo, Z., Wu, X., Chen, Y., 2022. Eikonal tomography with physics-informed neural networks: Rayleigh wave phase velocity in the northeastern margin of the tibetan plateau. Geophysical Research Letters 49, e2022GL099053. doi:10.1029/2022GL099053.

Cheng, G., Xie, X., Han, J., Guo, L., Xia, G.S., 2020. Remote sensing image scene classification meets deep learning: Challenges, methods, benchmarks, and opportunities. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 13, 3735–3756. doi:10.1109/JSTARS.2020.3005403.

Cheng, Z., Wu, Y., Li, Y., Cai, L., Ihnaini, B., 2025. A comprehensive review of explainable artificial intelligence (XAI) in computer vision. Sensors 25, 4166. doi:10.3390/s25134166.

Cuomo, S., Schiano Di Cola, V., Giampaolo, F., Rozza, G., Raissi, M., Piccialli, F., 2022.

Scientific machine learning through physics-informed neural networks: Where we are and what's next. Journal of Scientific Computing 92, 88. doi:10.1007/s10915-022-01939-z.

Cybenko, G., 1989. Approximation by superpositions of a sigmoidal function. Mathematics of Control, Signals and Systems 2, 303–314. doi:10.1007/BF02551274.

Doshi-Velez, F., Kim, B., 2017. Towards a rigorous science of interpretable machine learning. URL: https://arxiv.org/abs/1702.08608, arXiv:1702.08608.

Farhadpour, S., Warner, T.A., Maxwell, A.E., 2024. Selecting and interpreting multiclass loss and accuracy assessment metrics for classifications with class imbalance: Guidance and best practices. Remote Sensing 16, 533. doi:10.3390/rs16030533.

Ganotra, R., Gupta, S., Dora, S., 2024. A joint autoencoder and classifier deep neural network for ad and mci classification. International Journal of Imaging Systems and Technology 34, e23054. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/ima.23054, doi:https://doi.org/10.1002/ima.23054, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/ima.23054.

Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. MIT Press. http://www.deeplearningbook.org.

Haghighat, E., Juanes, R., 2021. Sciann: A keras/tensorflow wrapper for scientific computations and physics-informed deep learning using artificial neural networks. Computer Methods in Applied Mechanics and Engineering 373, 113552. doi:10.1016/j.cma.2020.113552. arXiv:2005.08803 [cs].

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778. doi:10.1109/CVPR.2016.90.

Hodson, T.O., 2022. Root-mean-square error (RMSE) or mean absolute error (MAE): When to use them or not. Geoscientific Model Development 15, 5481–5487. doi:10.5194/gmd-15-5481-2022.

Höhl, A., Obadic, I., Fernández-Torres, M.Á., Najjar, H., Oliveira, D.A.B., Akata, Z., Dengel, A., Zhu, X.X., 2024. Opening the black box: A systematic review on explainable artificial intelligence in remote sensing. IEEE Geoscience and Remote Sensing Magazine 12, 261–304. doi:10.1109/MGRS.2024.3467001.

Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. Neural Networks 2, 359–366. doi:10.1016/0893-6080(89)90020-8.

Huber, P.J., 1964. Robust estimation of a location parameter. The Annals of Mathematical Statistics 35, 73–101. doi:10.1214/aoms/1177703732.

Jordan, Mitchell, 2015. Machine learning: Trends, perspectives, and prospects. Science 349, 255–260. doi:10.1126/science.aaa8415.

Kamakshi, V., Krishnan, N.C., 2023. Explainable image classification: The journey so far and the road ahead. AI 4, 620–651. doi:10.3390/ai4030033.

Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L., 2021. Physics-informed machine learning. Nature Reviews Physics 3, 422–440. doi:10.1038/s42254-021-00314-5.

Kaufman, S., Rosset, S., Perlich, C., Stitelman, O., 2012. Leakage in data mining: Formulation, detection, and avoidance. ACM Transactions on Knowledge Discovery from Data 6, 1–21. doi:10.1145/2382577.2382579.

Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization URL: http://arxiv.org/abs/1412.6980, doi:10.48550/arXiv.1412.6980. arXiv:1412.6980 [cs].

Klivans, A., Stavropoulos, K., Vasilyan, A., 2024. Testable learning with distribution shift, in:

Proceedings of the 37th Annual Conference on Learning Theory.

Kozmin, A., Kalashev, O., Chernenko, A., Redyuk, A., 2025. Semi-supervised learned autoencoder for classification of events in distributed fibre acoustic sensors. Sensors 25, 3730. doi:10.3390/s25123730.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet classification with deep convolutional neural networks, in: Advances in Neural Information Processing Systems, pp. 1097–1105.

Krueger, D., Caballero, E., Jacobsen, J.H., Zhang, A., Binas, J., Zhang, D., Le Priol, R., Courville, A., 2021. Out-of-distribution generalization via risk extrapolation, in: Proceedings of the 38th International Conference on Machine Learning.

Kuznietsov, A., Gyevnar, B., Wang, C., Peters, S., Albrecht, S.V., 2024. Explainable AI for safe and trustworthy autonomous driving: A systematic review. IEEE Transactions on Intelligent Transportation Systems 25, 19342–19364. doi:10.1109/TITS.2024.3474469.

LeCun, Y., Bengio, Y., 1998. Convolutional networks for images, speech, and time series, in: Arbib, M.A. (Ed.), The Handbook of Brain Theory and Neural Networks. MIT Press, Cambridge, Massachusetts, pp. 255–258.

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. Nature 521, 436–444. doi:10.1038/nature14539.

Li, J., Lin, Y., Yao, H., Xu, J., Zhang, K., 2025. A core-surface flow model derived from magnetic data of macau science satellite-1. Earth and Planetary Physics 9, 577–585. doi:10.26464/epp2025034.

Li, J., Lin, Y., Zhang, K., 2024. Dynamic mode decomposition of the core surface flow inverted from geomagnetic field models. Geophysical Research Letters 51, e2023GL106362. URL: https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2023GL106362, doi:10.1029/2023GL106362.

Li, P., Pei, Y., Li, J., 2023. A comprehensive survey on design and application of autoencoder in deep learning. Applied Soft Computing 138, 110176. doi:10.1016/j.asoc.2023.110176.

Linardatos, P., Papastefanopoulos, V., Kotsiantis, S., 2021. Explainable ai: A review of machine learning interpretability methods. Entropy 23. URL: https://www.mdpi.com/1099-4300/23/1/18, doi:10.3390/e23010018.

Livermore, P.W., Wu, L., Chen, L., de Ridder, S., 2024. Reconstructions of jupiter's magnetic field using physics-informed neural networks. Monthly Notices of the Royal Astronomical Society 533, 4058–4067. doi:10.1093/mnras/stae1928.

Loshchilov, I., Hutter, F., 2017. Fixing weight decay regularization in adam. CoRR abs/1711.05101. URL: http://arxiv.org/abs/1711.05101, arXiv:1711.05101.

Lu, Z., Pu, H., Wang, F., Hu, Z., Wang, L., 2017. The expressive power of neural networks: A view from the width, in: Advances in Neural Information Processing Systems.

Lundberg, S.M., Lee, S.I., 2017. A unified approach to interpreting model predictions, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, Curran Associates Inc., Red Hook, NY, USA. p. 4768–4777.

Mao, A., Mohri, M., Zhong, Y., 2023. Cross-entropy loss functions: Theoretical analysis and applications. URL: https://arxiv.org/abs/2304.07288, arXiv:2304.07288.

Mao, J., Guo, S., Yin, X., Chang, Y., Nie, B., Wang, Y., 2025. Medical supervised masked autoencoder: Crafting a better masking strategy and efficient fine-tuning schedule for medical image classification. Applied Soft Computing 169, 112536. URL: https://www.sciencedirect.com/science/article/pii/S1568494624013103, doi:https://doi.org/10.1016/j.asoc.2024.112536.

Masci, J., Meier, U., Cireşan, D., Schmidhuber, J., 2011. Stacked convolutional auto-encoders for hierarchical feature extraction, in: Artificial Neural Networks and Machine Learning, Springer. pp. 52–59. doi:10.1007/978-3-642-21735-7_7.

Mehta, P., Bukov, M., Wang, C.H., Day, A.G.R., Richardson, C., Fisher, C.K., Schwab, D.J., 2019. A high-bias, low-variance introduction to machine learning for physicists. Physics Reports 810, 1–124. doi:10.1016/j.physrep.2019.03.001.

Muhammad, D., Bendechache, M., 2024. Unveiling the black box: A systematic review of explainable artificial intelligence in medical image analysis. Computational and Structural Biotechnology Journal 24, 542–560. doi:10.1016/j.csbj.2024.08.005.

Raghu, M., Poole, B., Kleinberg, J., Ganguli, S., Sohl-Dickstein, J., 2017. On the expressive power of deep neural networks, in: Proceedings of the 34th International Conference on Machine Learning, pp. 2847–2854.

Raissi, Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378, 686–707. doi:10.1016/j.jcp.2018.10.045.

Raissi, Yazdani, A., Karniadakis, G.E., 2020. Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. Science (New York, N.Y.) 367, 1026–1030. doi:10.1126/science.aaw4741.

Ramachandran, P., Zoph, B., Le, Q.V., 2017. Swish: a self-gated activation function URL: http://arxiv.org/abs/1710.05941. arXiv:1710.05941 [cs].

Ribeiro, M.T., Singh, S., Guestrin, C., 2016. "why should i trust you?": Explaining the predictions of any classifier. URL: https://arxiv.org/abs/1602.04938, arXiv:1602.04938.

Saeedi, J., Giusti, A., 2023. Semi-supervised visual anomaly detection based on convolutional autoencoder and transfer learning. Machine Learning with Applications 11, 100451. URL: https://www.sciencedirect.com/science/article/pii/S266682702300004X, doi:https://doi.org/10.1016/j.mlwa.2023.100451.

Schiassi, E., De Florio, M., Ganapol, B.D., Picca, P., Furfaro, R., 2022. Physics-informed neural networks for the point kinetics equations for nuclear reactor dynamics. Annals of Nuclear Energy 167, 108833. doi:10.1016/j.anucene.2021.108833.

Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization, in: Proceedings of the IEEE International Conference on Computer Vision (ICCV).

Sharma, P., Chung, W.T., Akoush, B., Ihme, M., 2023. A review of physics-informed machine learning in fluid mechanics. Energies 16, 2343. doi:10.3390/en16052343.

Simonyan, K., Vedaldi, A., Zisserman, A., 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps, in: International Conference on Learning Representations. URL: https://arxiv.org/abs/1312.6034, doi:10.48550/arXiv.1312.6034.

Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M., 2017. Smoothgrad: removing noise by adding noise. URL: https://arxiv.org/abs/1706.03825, arXiv:1706.03825.

Snoek, J., Larochelle, H., Adams, R.P., 2012. Practical bayesian optimization of machine learning algorithms URL: http://arxiv.org/abs/1206.2944, doi:10.48550/arXiv.1206.2944. arXiv:1206.2944 [cs, stat].

Sola, J., Sevilla, J., 1997. Importance of input data normalization for the application of neural networks to complex industrial problems. IEEE Transactions on Nuclear Science 44, 1464–1468. doi:10.1109/23.589532.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning

Research , 1929–1958.

Su, S.Z., Liu, Z.H., Xu, S.P., Li, S.Z., Ji, R., 2015. Sparse auto-encoder based feature learning for human body detection in depth image. Signal Processing 112, 43–52. URL: https://www.sciencedirect.com/science/article/pii/S0165168414005052, doi:https://doi.org/10.1016/j.sigpro.2014.11.003. signal Processing and Learning Methods for 3D Semantic Analysis.

Varoquaux, G., Colliot, O., 2023. Evaluating machine learning models and their diagnostic value, in: Colliot, O. (Ed.), Machine Learning for Brain Disorders. Humana, New York, NY. chapter 20. doi:10.1007/978-1-0716-3195-9_20.

Vyas, N., Morwani, D., Zhao, R., Kwun, M., Shapira, I., Brandfonbrener, D., Janson, L., Kakade, S., 2025. Soap: Improving and stabilizing shampoo using adam. URL: https://arxiv.org/abs/2409.11321, arXiv:2409.11321.

Wagenmaker, A., Huang, K., Ke, L., Jamieson, K., Gupta, A., 2024. Overcoming the sim-to-real gap: Leveraging simulation to learn to explore for real-world reinforcement learning, in: Advances in Neural Information Processing Systems.

Xing, C., Ma, L., Yang, X., 2016. Stacked denoise autoencoder based feature extraction and classification for hyperspectral images. Journal of Sensors 2016, 3632943. URL: https://onlinelibrary.wiley.com/doi/abs/10.1155/2016/3632943, doi:https://doi.org/10.1155/2016/3632943, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1155/2016/3632943.

# Chapter 3

# Research Questions

This thesis investigates how these advances in machine learning can be applied to the study of both fluid flow at the top of the Earth's outer core and thermal conditions at the core–mantle boundary. In doing so, it is hoped that the following research questions will be answered.

- **Can regional core-flow inversions recover the large-scale structures found in global flow models?**

  Core-surface flow is commonly inferred using global spherical harmonic representations, which have identified an eccentric, anticyclonic gyre beneath the Atlantic hemisphere, an accelerating high-latitude jet beneath the Bering Strait, and changes in equatorial azimuthal flow beneath Indonesia and Central America. However, the global basis and large-scale regularisation may blur separate regional flow features into a single planetary-scale circulation. The gyre may therefore not be recovered when the flow is constrained only by regional magnetic field and secular variation data. Previous regional inversions have had mixed success and remain sensitive to smoothing or prior assumptions (Rogers, 2022; Schwaiger et al., 2023). Physics-informed neural networks provide an alternative approach, by recovering regional flows that fit the local secular variation while satisfying an imposed physical constraint. A mosaic of independently inferred flows can therefore test whether the gyre connects across regional boundaries, and whether it is supported by the magnetic-field data rather than the global modelling approach.

- **Do localised flow models contain information helpful for forecasting the Secular Variation of the Earth's magnetic field?**

  Secular variation forecasts commonly use global field models, statistical extrapolation, data assimilation or steady core-flow advection (Alken et al., 2021; Finlay et al., 2020; Fournier et al., 2021; Brown et al., 2021). However, global flow representations may smooth out localised structures that contain useful information about future field changes. The regional PINN methodology developed to address the first research question provides a way to test this by using a set of independently inferred local flows within a steady-flow advection scheme. Comparing hindcasts produced using regional and global flows can therefore determine whether the additional small-scale information retained by localised models improves secular variation forecasts.

- **Can heat flux heterogeneity be inferred from the radial magnetic field at the Core Mantle Boundary?**

  CMB heat-flux heterogeneity is commonly estimated from seismic tomography, but its amplitude and structure remain poorly constrained because seismic velocity variations have both thermal and compositional origins (Masters et al., 1996; Zhao et al., 2015). Numerical geodynamo simulations show that heterogeneous heat flux can influence the morphology of the radial magnetic field, although similar features may also arise from other model parameters or natural magnetic variability. A convolutional autoencoder classifier provides a way to test whether magnetic-field snapshots contain sufficient information to distinguish between different classes of imposed heat-flux heterogeneity. Applying this approach to controlled and more varied simulation sets, followed by models of the Earth's magnetic field, can test how reliably heat-flux heterogeneity can be recovered and identify the magnetic-field features most important for classification.

# References

Alken, P., Chulliat, A., Nair, M., 2021. Noaa/ncei and university of colorado candidate models for igrf-13. Earth, Planets and Space 73, 44. URL: https://doi.org/10.1186/s40623-020-01313-z, doi:10.1186/s40623-020-01313-z.

Brown, W.J., Beggan, C.D., Cox, G.A., Macmillan, S., 2021. The bgs candidate models for igrf-13 with a retrospective analysis of igrf-12 secular variation forecasts. Earth, Planets and Space 73, 42. URL: https://doi.org/10.1186/s40623-020-01301-3, doi:10.1186/s40623-020-01301-3.

Finlay, C.C., Kloss, C., Olsen, N., Hammer, M.D., Tøffner-Clausen, L., Grayver, A., Kuvshi-

nov, A., 2020. The chaos-7 geomagnetic field model and observed changes in the south atlantic anomaly. Earth, Planets and Space 72, 156. URL: https://doi.org/10.1186/s40623-020-01252-9, doi:10.1186/s40623-020-01252-9.

Fournier, A., Aubert, J., Lesur, V., Ropp, G., 2021. A secular variation candidate model for igrf-13 based on swarm data and ensemble inverse geodynamo modelling. Earth, Planets and Space 73, 43. URL: https://doi.org/10.1186/s40623-020-01309-9, doi:10.1186/s40623-020-01309-9.

Masters, T.G., Johnson, S., Laske, G., Bolton, H., 1996. A shear - velocity model of the mantle. Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 354, 1385–1411. doi:10.1098/rsta.1996.0054.

Rogers, H.F., 2022. Improving inversions of regional outer core surface flow models through the application of spherical slepian function URL: https://era.ed.ac.uk/handle/1842/39295, doi:10.7488/era/2546. accepted: 2022-08-02T10:34:07Z.

Schwaiger, T., Jault, D., Gillet, N., Schaeffer, N., Mandea, M., 2023. Local estimation of quasi-geostrophic flows in earth's core. Geophysical Journal International 234, 494–511. URL: https://doi.org/10.1093/gji/ggad089, doi:10.1093/gji/ggad089, arXiv:https://academic.oup.com/gji/article-pdf/234/1/494/50509990/ggad089.pdf.

Zhao, C., Garnero, E.J., McNamara, A.K., Schmerr, N., Carlson, R.W., 2015. Seismic evidence for a chemically distinct thermochemical reservoir in earth's deep mantle beneath hawaii. Earth and Planetary Science Letters 426, 143–153. URL: https://www.sciencedirect.com/science/article/pii/S0012821X15003684, doi:https://doi.org/10.1016/j.epsl.2015.06.012.

# Chapter 4

# Local Flow estimation at the top of the Earth's core using Physics Informed Neural Networks

Naomi Shakespeare-Rees, Philip W. Livermore, Christopher J. Davies, Hannah F. Rogers, William J. Brown, Ciarán D. Beggan, Christopher C. Finlay

Published in *Physics of the Earth and Planetary Interiors* as Shakespeare-Rees et al. (2025). Minor editorial changes have been made for consistency with the thesis.

## Abstract

The Earth's main geomagnetic field arises from the continuous motion of the fluid outer core. By assuming that the field changes are advection-dominated, and that diffusion only plays a minor role, the fluid motion at the core surface can be related to the secular variation of the geomagnetic field, providing an observational approach to understanding the motions in the deep Earth. The majority of existing core flow models are global, showing features such as an eccentric planetary gyre, with some evidence of rapid regional changes. By construction, the flow defined at any location by such a model depends on all magnetic field variations across the entire core-mantle boundary: because of this nonlocal dependence of the flow on the magnetic field, it is very challenging to interpret local structures in the flow as due to specific local changes in magnetic field. Here we present an alternative strategy in which we construct regional flow models that rely only on local secular changes. We use a novel technique based on machine learning termed Physics-Informed Neural Networks (PINNs), in which we seek a regional flow model that simultaneously fits both the local magnetic field variation and dynamical conditions assumed satisfied by the flow. Although we present results using the Tangentially Geostrophic flow constraint, we set out a modelling framework for which the physics constraint can be easily changed by altering a single line of code. After validating the PINN-based method on synthetic flows, we apply our method to the CHAOS-8.1 geomagnetic field model, itself based on data from Swarm. Constructing a global mosaic of regional flows, we reproduce the planetary gyre, providing independent evidence that the strong secular changes at high latitude and in equatorial regions are part of the same global feature. Our models also corroborate regional changes in core flows over the last decade. In our models, we find that the azimuthal flow under South America has changed sign quasi-periodically, with a recent sign change in 2022. Furthermore, our models endorse the existence of a dynamic high latitude jet, which began accelerating around 2005 but has been weakening since 2017.

## 4.1 Introduction

The Earth's main magnetic field, which is generated by a self-sustaining geodynamo arising from fluid motions in the Earth's core (Bullard, 1950), exhibits fluctuations on timescales of years to millennia and longer (Constable and Constable, 2023). Changes in the geomagnetic field on timescales of years to millennia are termed Secular Variation (SV) (Jackson and Finlay, 2015). The geomagnetic field has been measured by networks of ground-based observatories since 1837 (Macmillan, 2007), and supplemented by continuous satellite measurements since 1999 (Jackson and Finlay, 2015; Friis-Christensen et al., 2006; Olsen and Floberghagen, 2018; Zhang, 2023; Jiang et al., 2024). Changes in the geomagnetic field can have critical impacts on both industry and scientific exploration in a diverse range of disciplines, such as navigation, satellite operations and the protection of the Earth from space weather. By mapping the outer core flow that generates the SV, the dynamics and properties of the core can be explored, and SV forecasts such as the candidate models of IGRF-13 (Alken et al., 2021) constructed.

The typical approach to map core flows is to use observations from satellites, observatories and other surveys to construct global geomagnetic field models using spherical harmonics, which can be downward-continued to the Core-Mantle Boundary (CMB). The inversion of the SV to recover motions on the outer core is under-determined, and so additional assumptions must be included. A widely used assumption is that of *frozen flux* (Roberts and Scott, 1965), where on sufficiently short timescales, magnetic diffusion can be considered negligible. Under this simplification, the magnetic field lines are 'frozen' into packets of moving liquid at the surface of the outer core, so that the field becomes a tracer for the flow (Bloxham and Jackson, 1991). However, even with this assumption, the problem of non-uniqueness remains (Backus, 1968). There are still more unknown components of the flow than the number of equations describing them. This means additional constraints on the flow are required in order to reduce the fundamental non-uniqueness. Various flow assumptions have been used in previous studies, including steady flows (Gubbins et al., 1982), toroidal only flow (Whaler, 1980), tangential geostrophy (TG) (Le Mouël, 1984; Rau et al., 2000), quasi-geostrophic flow (Pais and Jault, 2008) and helical and columnar flow (Amit and Olson, 2004). Non-uniqueness is further reduced by the large scale flow assumption, in which small scale flow structures are penalised (e.g. Bloxham, 1988).

Spherical harmonics are also used to create these global flow models, and underpin most studies to date. Results from previous global outer core flow inversions, as well as numerical simulations, reveal an eccentric, anticyclonic, planetary gyre (Pais and Jault, 2008; Gillet et al., 2022; Kloss and Finlay, 2019). This gyre is a planetary-scale circulation pattern, travelling west under the Atlantic, north west under North America, westwards near the North Pole, then south under Asia. Although this is a global pattern, there are several regional features that are of specific interest. Firstly, there is evidence for a localised jet under the Bering strait, which has been strengthening since 2005 and is associated with a strong change in the SV at high northern latitudes (Livermore et al., 2017). Additionally, multiple studies, such as Whaler et al. (2022) and Li et al. (2024), have observed changes in the azimuthal flow direction in areas in the equatorial region, particularly beneath Indonesia and Central America. These may be a result of hydromagnetic waves producing changes on interannual timescales (Gillet et al., 2022), and many of these sign changes may be associated with geomagnetic jerks - a phenomenon in which the SV changes rapidly, sometimes in a spatially localised manner (Brown et al., 2013).

In contrast to global core flow inversions, we adopt a local approach in which flows are inferred from a regional realisation of a global geomagnetic field model, presenting an independent method to check global flow analysis. This is important, as a large-scale global inversion will act to interpolate multiple local features into one larger one. For example, a global approximation to westward drift in the Atlantic and flow into the polar regions would be a gyre that connects these regions, and so the use of regional inversions could test how robust these features are. Additionally, in a global flow analysis, any point on the CMB depends non-locally on all the points in the model, but in a local flow inversion the same point would only depend on the other points in the region. As the data coverage of the Earth is spatially uneven, especially at the poles, a regional methodology would probe the reliability of local features in these areas, ensuring they are not an artefact of a global inversion. This would allow for regional features to be studied in more detail, without the uncertainty of coupling to neighbouring regions.

Local core flow inversions have been attempted in two previous studies with mixed results. Rogers (2022) studied regional variations using spherical Slepian functions. This technique produced better separation of SV at the Earth's surface compared to spherical harmonics but reliable local flow separations and SV separations at the CMB were not achieved. Schwaiger et al. (2023) presented a local core flow inversion methodology based on pointwise inversion,

which was able to reproduce the main features found in global core flow studies, but it was found that additional smoothing was required to prevent unreliable re-construction of small-scale flows. They also found that their results heavily depended on what prior they used to reduce the non-uniqueness.

This work introduces a novel approach to infer local flows, employing recent advancements in machine learning through the use of Physics-Informed Neural Networks (PINNs). Machine learning refers to a set of statistical techniques to leverage data in order to undertake a task, without being explicitly programmed to do so (Armstrong and Fletcher, 2019). This allows the user to extract knowledge and draw inferences from data (Jordan and Mitchell, 2015). Neural Networks (NN) are a machine learning technique, consisting of layers of artificial neurons that can process information. PINNs, first proposed in Raissi et al. (2019), are a class of neural network that use mathematical descriptions of physical laws as constraints in order to solve forward and inverse problems. PINNs have been used in a diverse range of fields such as fluid mechanics (Raissi et al., 2020), medicine (Arzani et al., 2021), nuclear physics (Schiassi et al., 2022) and seismology (Chen et al., 2022). PINNs have been recently used in core flow and length of day (LOD) analysis by Li et al. (2024), albeit in a global flow inversion framework, rather than the local approach that we adopt in this study.

The methodology presented in this work takes secular variation from regional latitude-longitude boxes from the global geomagnetic field model CHAOS-8.1 as input, and then outputs a flow that both re-produces the input secular variation and satisfies an additional flow constraint. We aim to establish a framework that can be used for any flow assumption, which we illustrate here using the TG flow constraint. We do not rely on any prior information other than the physics constraint. The methodology and validation is described in section 4.2. Results, including local flow analysis as well as a global flow model constructed from a mosaic of regional models, are presented in section 4.3. All of these results are discussed in section 4.4, with conclusion in section 4.5.

## 4.2 Method

### 4.2.1 Mathematical Framework

We model the Earth's core as a sphere, described using spherical coordinates $(r, \theta, \phi)$, and assume that the flow within it obeys a non-penetration boundary condition at the edge of the core such that $u_r = 0$ at the CMB. Assuming frozen flux, the radial magnetic field, $B_r$, can be related to the outer core surface horizontal fluid movement via the radial component of the induction equation, at the core surface, just below the CMB. This is written as

$$\frac{\partial B_r}{\partial t} = -\nabla_H \cdot (\boldsymbol{u} B_r), \tag{4.1}$$

in which $\boldsymbol{u} = [0, u_\theta, u_\phi]$ is the flow, which is sought in the inversion methodology, and $\nabla_H$ is the horizontal gradient operator.

The tangential geostrophic flow assumption arises from considering the force balance at the top of the Earth's core (Le Mouël, 1984). If the force balance is dominated by the Coriolis (rotational) and pressure forces, then the Navier-Stokes equation for the motion of the fluid reduces to

$$2\rho(\boldsymbol{\Omega} \times \boldsymbol{u})_H + \nabla_H p = 0, \tag{4.2}$$

where $\rho$ is the density of the liquid outer core, $\boldsymbol{\Omega}$ is the rotation vector of the Earth, and $p$ is the pressure of the fluid.

Taking the horizontal divergence of the cross product of equation (4.2) and the radial unit vector $\hat{\boldsymbol{r}}$ gives

$$\nabla_H \cdot (\boldsymbol{u} \cos\theta) = 0, \tag{4.3}$$

which can be written as

$$\frac{\tan\theta}{r} u_\theta - \nabla_H \cdot \boldsymbol{u} = 0, \tag{4.4}$$

that defines the TG constraint (Le Mouël (1984), Amit and Pais (2013)). This additional constraint modifies equation (4.1) to

$$\frac{\partial B_r}{\partial t} = -\boldsymbol{u} \cos\theta \cdot \nabla_H \left( \frac{B_r}{\cos\theta} \right). \tag{4.5}$$

It was noted by Backus and Le Mouël (1986) that this flow is non-unique on contours of $B_r/\cos\theta$, as well as in regions bounded by contours of $B_r/\cos\theta$ that do not cross the equator, meaning care must be taken when interpreting flows in these areas.

In spherical coordinates, the total flow $\mathbf{u}$ can be decomposed into toroidal and poloidal components

$$\boldsymbol{u} = \nabla \times T(\theta,\phi)\mathbf{r} + \nabla_H(rP(\theta,\phi)), \tag{4.6}$$

whose horizontal components can be written (Holme, 2015),

$$\boldsymbol{u}_T = \left(\frac{1}{\sin\theta}\frac{\partial T}{\partial \phi}, -\frac{\partial T}{\partial \theta}\right), \tag{4.7}$$

$$\boldsymbol{u}_P = \left(\frac{\partial P}{\partial \theta}, \frac{1}{\sin\theta}\frac{\partial P}{\partial \phi}\right). \tag{4.8}$$

Therefore, in our flow inversion we seek toroidal ($T$) and poloidal ($P$) scalar functions which define the flow.

### 4.2.2 Neural Networks

Our inversion methodology consists of two Fully Connected Neural Networks (FC-NNs) working in parallel: one to describe the toroidal scalar $T$ and the other to describe the poloidal scalar $P$. A schematic for this is shown in figure 4.1. FC-NNs consist of layers of nodes, organised into an input layer, multiple hidden layers, and an output layer (Jordan and Mitchell, 2015). Each of these nodes are connected to each other and have an associated weight and bias, and the output of each node is computed by some non-linear function of the input and the weight. This non-linear function is called the activation function, and the weights and biases are adjusted during a process termed training (Goodfellow et al., 2016). Once the network weights are determined, the flows are then described using equations (A.1) and (4.8).

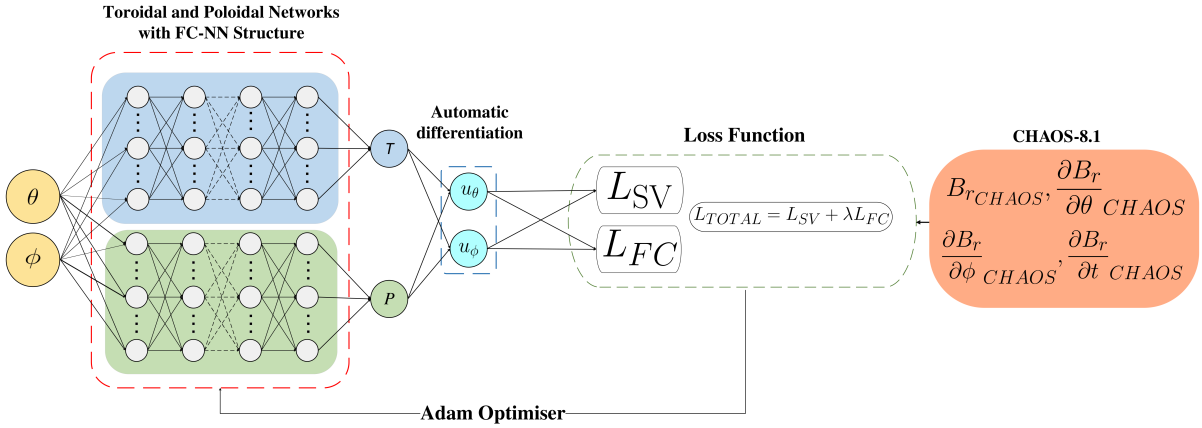



Figure 4.1: Schematic of the PINN used in this study, where $\theta, \phi$ are co-latitude and longitude, $T$ and $P$ are the toroidal and poloidal values on the $\theta, \phi$ grid, and $u_\theta$ and $u_\phi$ are the components of the horizontal flow. Yellow indicates inputs to the network, the two parallel FC-NNs are blue and green, loss functions in white, inputs to the loss function in orange, and the outputs are cyan.

While we could train a network without any flow assumptions, here we use an additional physics constraint to enforce TG. This takes place solely in the training stage, which is done by minimising a loss term consisting of a data loss, which aims to implement equation (4.1), and a physics-based loss, which aims to implement equation (4.4). In this way, we fit both to the data and to the underlying physics. This is so we can establish a framework wherein different flow assumptions, such as Toroidal or Helical flow, can be swapped in and out using a single line of code. All the loss terms are implemented by defining a derived quantity from the network and a corresponding target, with the loss value determined by the Mean Squared Error (MSE)

between the quantity and the target. The total loss, $L_{TOTAL}$ is defined by

$$L_{\mathrm{TOTAL}} = L_{\mathrm{SV}} + \lambda L_{\mathrm{FC}}, \tag{4.9}$$

which consists of two loss terms:

- The Data Loss ($L_{\mathrm{SV}}$):

$$L_{\mathrm{SV}} = \frac{1}{N}\sum_i \left[ \frac{\partial B_r}{\partial t}_{PINN} - \frac{\partial B_r}{\partial t}_{CHAOS} \right]^2 \Bigg|_{\theta_i,\phi_i} \tag{4.10}$$

  where $N$ is the number of points, and the SV calculated from the horizontal flows is

$$\frac{\partial B_r}{\partial t}_{PINN} = -\frac{1}{r}\left( \frac{u_\phi}{\sin\theta} \frac{\partial B_r}{\partial \phi}_{CHAOS} + u_\theta \frac{\partial B_r}{\partial \theta}_{CHAOS} \right) - B_{r_{CHAOS}} (\nabla_H \cdot \boldsymbol{u}), \tag{4.11}$$

- and the TG Flow Loss ($L_{\mathrm{FC}}$):

$$L_{\mathrm{FC}} = \frac{1}{N}\sum_i \left[ \nabla_H \cdot \boldsymbol{u} - \left( \frac{\tan\theta}{r} \right) u_\theta \right]^2 \Bigg|_{\theta_i,\phi_i} \tag{4.12}$$

where $\theta_i, \phi_i$ are the pre-determined set of latitude-longitude grid points on which we impose the constraints. We use the same grid to constrain the data and the physics. For the data constraint, the target is the SV from CHAOS-8.1, whereas in the flow constraint the target is zero. A scaling analysis of typical magnitudes of the two terms suggests a value of $\lambda = 1000$ for the weighting factor, which is applied to $L_{FC}$ in equation (4.9) so that both terms are $\mathcal{O}(1)$. This is desirable so that one loss term is not more important than the other.

The networks are trained using the optimisation method Adam (Kingma and Ba, 2017), which updates both networks simultaneously. The PINN takes in a $(\theta, \phi)$ grid of points, at a radius of 3485 km, and then analytically computes the toroidal and poloidal scalar value at each of those points, as well as the derivatives. These derivatives are then used in equations (A.1) and (4.8) to produce $u_\theta, u_\phi$. The total error is calculated by summing $L_{SV}$, the error between the network derived SV and the CHAOS-8.1 SV, and $L_{FC}$, the departure of the flow from the tangentially geostrophic flow assumption, together with a weighting factor $\lambda$. This error is then back-propagated through the network by the Adam optimiser, adjusting the weights and biases of the network to attempt to minimise the MSE error. This process repeats for 100,000

iterations until the loss is minimised such that the loss does not decrease further. The PINN methodology is written with PyTorch Version 2.5 (Paszke et al., 2019). Our dataset is of order 1000-10000 points, which is modest by machine learning standards, and so we are able to use all the data in each iteration of training. When training the network, we initialise the weights from a pseudo-random distribution using a technique called Xavier initialisation (Glorot and Bengio, 2010). The optimiser then implements a descent method to seek a global minimum. However, we found empirically that the loss function is very complex, and so the optimiser will almost inevitably only find a local minima. We therefore choose multiple different seeds – and so initial weights – for independent training, from which we select the model that achieves the lowest loss.

### 4.2.3 Training Dataset: CHAOS-8.1

Our training procedure requires values for the SV to train the PINN, which enters through $L_{SV}$. We use the model CHAOS-8.1 (Kloss et al., 2024) projected onto a grid in spherical coordinates. CHAOS-8.1 spans from 1999 to 2025, and is built from satellite and ground based observatory data. It is the latest update of the CHAOS family of models, which have been used for multiple core flow inversion studies, such as Gillet et al. (2022), making it a suitable choice of model for this study. The time-dependent internal field model is defined up to spherical harmonic degree 20, but only the coefficients for the main field and SV up to degree 13 are used here, as at higher degrees the crustal field dominates the magnetic field signal (Langel and Estes, 1982). CHAOS-8.1 is also temporally regularised in order to reduce non-uniqueness, by taking temporal covariances from geodynamo simulations. This is a departure from the regularisation method of CHAOS-7 (Finlay et al., 2020), which penalised the second time derivative at the endpoints, and the third time derivative throughout. This difference in regularisation method affects the small scale features of the SV and the intensity of the acceleration, which may have an effect on the recovered flows (Kloss et al., 2024).

In order to use CHAOS-8.1 as input to the local flow inversion, a subsection of the SV is selected, either in a latitude-longitude box (as in figure 4.2), or in a longitudinal band, each with a grid density of one point per degree. A 5° border is added to the area of interest on all edges, which is removed after training. This is to promote continuity between adjacent regions, because continuity is not explicitly imposed, as well as to avoid edge effects. Tests with larger overlaps did not significantly change the results. To maximise the performance of the model,

and to minimise training time, the CHAOS-8.1 input and outputs of the network are re-scaled so that they all have a magnitude of about 1 (Sola and Sevilla, 1997). We do this by measuring $B_r$ in $\mu$T, time in 0.1 yr, and length in km. Typical values of $B_r$ at the CMB are 500 $\mu$T, and so typical values of SV are then 1 $\mu$T/0.1 yr in these units, while flow speeds are typically 1 km/0.1 yr.

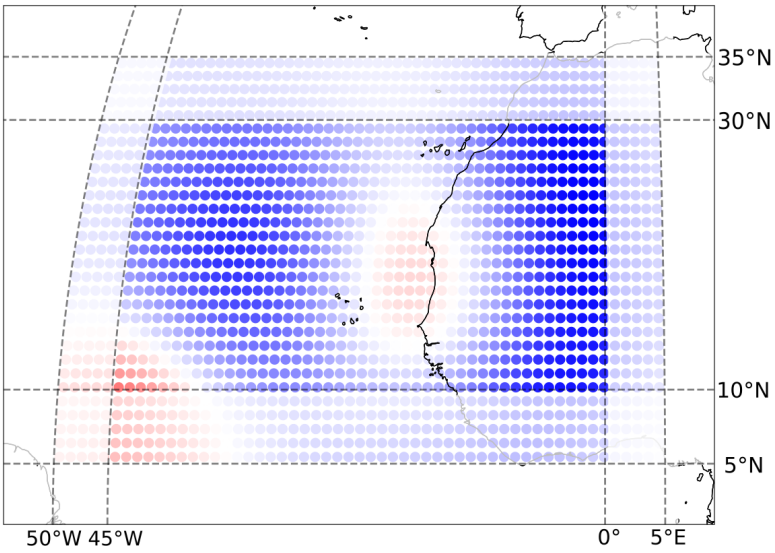



Figure 4.2: Example of a latitude-longitude box used for inversion. After training, a 5° border shown in translucent colouring, is removed. Colours show the radial SV at the CMB from CHAOS-8.1 for reference.

### 4.2.4 Parameter Choices

Having described the algorithm, there remain several parameter choices that define both the model and how it is trained. First, is the learning rate of the model, which governs how rapidly the weights change with each iteration. After testing a variety of learning rates, we found empirically that a standard learning rate of $10^{-3}$ worked well. Second is the choice of activation function. We found that the hyperbolic tangent function worked well, which is widely used due to its zero-centred property that results in faster training (Goodfellow et al., 2016). Lastly, we must select a network size - the number of hidden layers and the number of neurons per layer. A small network has a very limited functional representation, and in our case would result in a flow too smooth to fit the SV data. In contrast, a large network can represent spatially complex behaviour, and could result in flows that fit the SV but have spurious small-scale features. We seek to adopt an optimal network size that is large enough to fit the data constraints, but no larger, which effectively penalises small-scale features of the flow. This approach is analogous

to the spatial regularisation of global flows, for which their complexity is penalised through an explicit trade-off between data fit and complexity (for example, the strong norm presented in Bloxham (1988)). To be clear, we do not impose explicit regularisation of our flows, although we could by introducing an additional loss term in equation (4.9).

Our aim is to find a single network size that we can use in each region that we consider. To find the optimum network, we test a variety of network sizes for areas underneath the Atlantic (10°N to 20°N latitude, 45°W to 5°E longitude) and the South China Sea (10°N to 20°N latitude, 100°E to 130°E longitude), which represent end member behaviour. Under the Atlantic we expect simple westward flow (Holme, 2007), whereas under the South China Sea we expect complex, diverging flow (Whaler et al., 2022). For each tested network, we record the SV Root Mean Squared Error (RMSE) and the complexity as measured by the average of the squared second spatial derivative of the flow (Bloxham, 1988), evaluated on the SV data grid. We repeat this with 5 different seeds to investigate the spread in results. These are then plotted as trade-off curves in a manner similar to Hansen and O'Leary (1993).

The results for both regions are plotted in figure 4.3, which shows trade-off curves of complexity against RMSE. The grey points show the local minima for all the different seeds, and the black points indicate the seed for each network size which has the minimum RMSE, which we adopt as the closest model we have obtained to the global minimum. Simple networks are on the right, as they have a low complexity but a high SV RMSE. Complex networks are on the left, as they have a low SV RMSE but a high complexity. The network size with the highest complexity is shown in blue, whereas the one with the lowest complexity is shown in orange. Examples of the recovered flows and SV residuals at these points are shown in figure 4.4 for the area under the Atlantic, and figure 4.5 for the area under the South China Sea. For both of these figures, the size of network that recovers the most complex flows fits the data better than the size that recovers the least complex flows, but with the unwanted side effect of adding very rapid, small scale flows. This is due to the non-uniqueness present, both the large and small scale flows fit the SV adequately. A subjective way to choose the size of network is to choose the one at the 'knee' of this trade-off curve, and the values around this point are shown in green. The bottom panels of figure 4.3 shows the points at the knee in more detail. The behaviour is not always monotonic, although there is a general trend, perhaps due to the training getting stuck at a local minimum. For the area under the South China Sea, it was more difficult to find

complex models which fit to the data, as larger networks did not converge. We choose 8 layers of 40 neurons, indicated as 8[40] in figure 4.3, which was consistently at the knee across regions. While both 8[40] and 4[40] would be an appropriate pick for the size of model, we opted for the larger model to ensure that there would be enough free parameters to fit any SV structure. Cross-Validation (Goodfellow et al., 2016), a method often used for model selection, was not used in this case due to our small dataset, as Cross-Validation requires the dataset to be split during training.

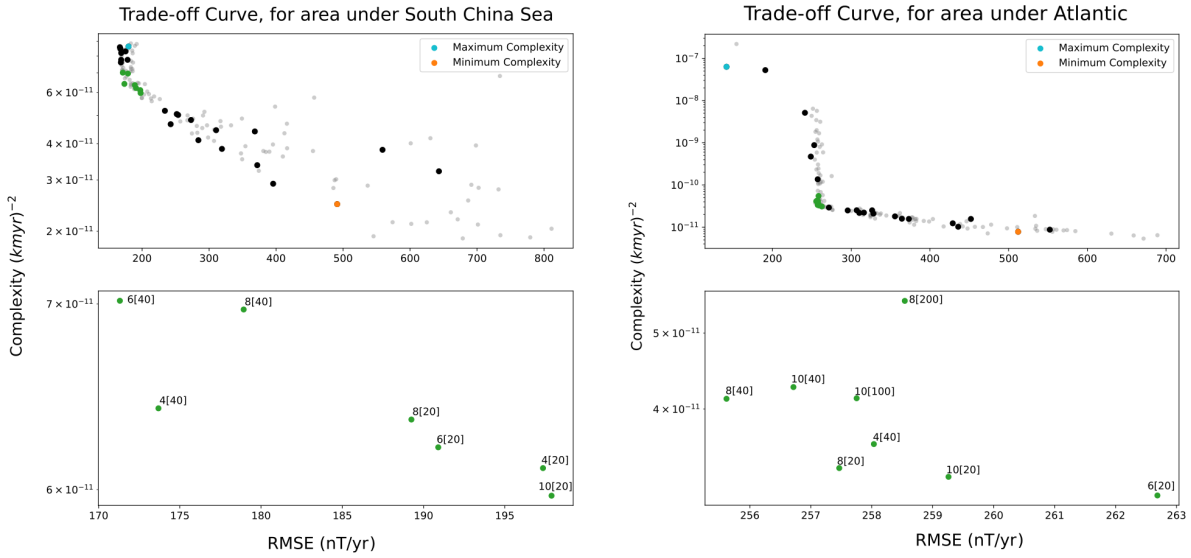



Figure 4.3: Trade-off curves for an area under the South China Sea (10°N to 20°N latitude, 100°E to 130°E longitude) (left) and Atlantic (10°N to 20°N latitude, 45°W to 5°E longitude) (right). Black dots indicate the minimum RMSE for each size of network, grey dots indicate the local minima, and the green dots indicate the points at the knee of the curve, which are shown in more detail on the bottom panel. The labels indicate the size of network, in the format $a[b]$, where $a$ is the number of layers and $b$ is the number of neurons per layer.

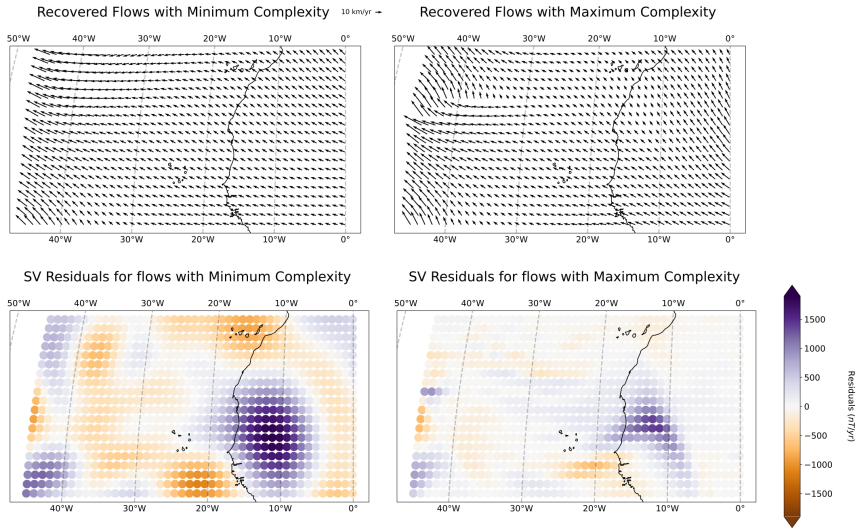



Figure 4.4: Recovered flows and SV Residuals for the network sizes that produce the least complex (left) and most complex (right) flows, in the area under the Atlantic, at 35°N to 5°N latitude and 0°W to 45°W longitude.

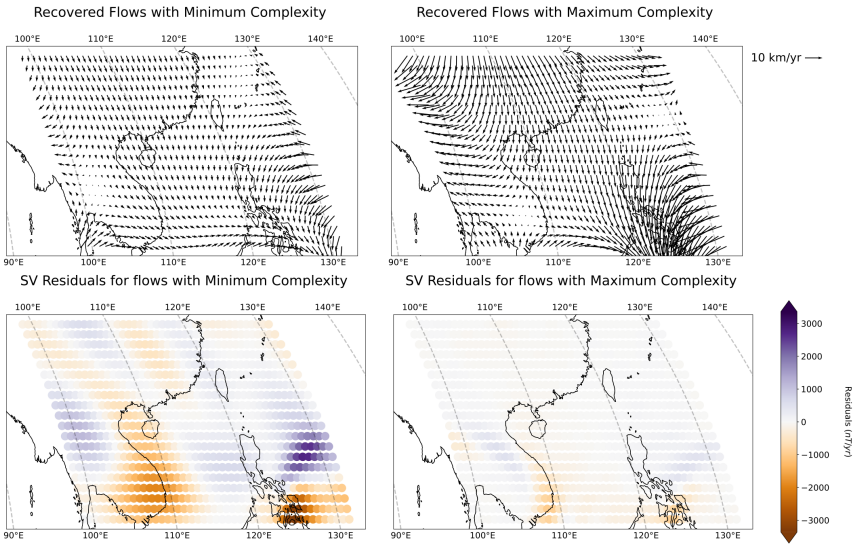



Figure 4.5: Recovered flows and SV Residuals for the network sizes that produce the least complex (left) and most complex (right) flows, in the area under the South China Sea.

In principle large networks will over-fit the data, which in this case would be evident through high complexities and spurious flows. We note that while the complexity varies across four orders of magnitude in the area under the Atlantic (shown in the left panels of figure 4.3), the complexity is never particularly high. This is even more apparent under the South China Sea (shown in the right panels of figure 4.3), where the complexity only varies over one order of magnitude. This may be a consequence of using CHAOS-8.1 for the input field and SV, as this model is already smoothed spatially and temporally, and we find that it is apparently difficult

to overfit the data in this case. Adding artificial Gaussian noise to the SV data increased the range of complexity magnitudes, and so perhaps smoothing through the choice of the size of networks will be more critical when using 'noisier', sparser or unevenly sampled data sources, such as Geomagnetic Virtual Observatories (GVOs) (Mandea and Olsen, 2006), rather than field models whose complexity was already controlled during their construction.

### 4.2.5 Synthetic Tests

In order to test the performance of the chosen networks, a number of synthetic examples were set up with pre-defined flows for the network to discover. On a latitude-longitude grid spanning -15° to 15° latitude and 0° to 55° longitude with 1-degree spacing in each direction, we define $B_r$ using

$$B_r = \exp\left(-\frac{(\theta - \frac{90\pi}{180})^2 + (\phi - \frac{30\pi}{180})^2}{2}\right), \tag{4.13}$$

which describes a single peak centred at (90°, 30°) in radians. For each flow, the instantaneous SV was determined using equation (4.1). A range of synthetic flows were investigated, as listed in table 4.1, which span behaviours from simple drift to a more complex vortex. The final synthetic flow tested was a rigid body rotation section of a circular Rankine vortex patch first presented in Rankine (1872), and most recently used in a geomagnetism context by Amit (2014). The description of this flow does not fit in the table, so it is described here. Rankine flow contains an inner circular area, described by a boundary of angular distance $H$, wherein the fluid motion follows that of solid body rotation with vorticity $\zeta$, and an outer area where the flow does not follow solid body rotation. We only use this inner section for our synthetic flow. The azimuthal toroidal flow, $u_t$, at a given angle from the centre, $h$, is given by

$$u_t = -\begin{cases} \frac{1}{2}\zeta h, & \text{if } h < H \\ \frac{1}{2}\zeta \frac{H^2}{h^2}, & \text{if } h > H, \end{cases} \tag{4.14}$$

where $H = 35°$ is the angular distance of the solid body circular area and $\zeta = 0.1km/0.1year$ is a constant vorticity in the region where $h < H$. The maximum flow speeds are $2.6km/0.1year$.

The models are trained with 5 different seeds, with a 5° border removed after training as in figure 4.2. We did not apply the TG flow restriction here, as a number of these flows do not satisfy it. Figure 4.6A shows the input SV, figure 4.6B shows the output SV from the PINN,

and figure 4.6C shows the residuals between the input and output SV, all for the rigid body rotation. Figure 4.6D shows the synthetic rigid body rotation flow used to generate the input SV and figure 4.6E shows the output flows from the PINN. Table 4.1 shows the description and results for other synthetic flows, with the error between the true value and the output value expressed using the relative L1 error given by:

$$Relative\ \%\ L1\ Error\ = \frac{\overline{(Value_{Predicted} - Value_{True})} \times 100}{\overline{Value_{True}}} \tag{4.15}$$

in which $\bar{}$ denotes the mean. The spread in the results due to the different starting seeds calculated by taking the standard deviations of the percentage errors are also shown. Each of these tests successfully recovers the synthetic flows, validating the method for large scale flows. The residuals also do not have any structures correlating to the original input, and the residuals are all small-scale.

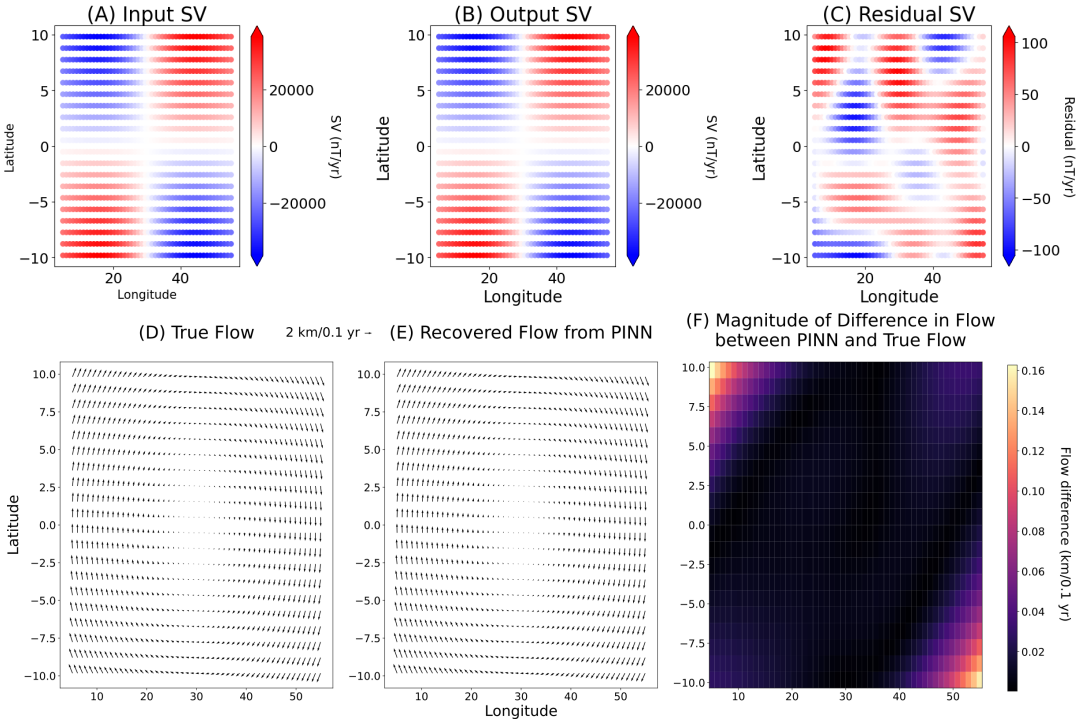



Figure 4.6: Example of the PINN recovery of the flow, based on a synthetic test using rigid body rotation (A) Input SV to the PINN, (B) Output SV from the PINN (C) Residual SV (Difference between input and output) (D) Synthetic flow used to generate input SV (E) Output flow from the PINN (F) Difference between output and synthetic flow. A 5° border has been removed after training as in figure 4.2.

| **Flow Description** | **Flow Equations** | **SV Relative L1 % Error** | $\boldsymbol{u_\theta}$ **Relative L1 % Error** | $\boldsymbol{u_\phi}$ **Relative L1 % Error** |
|---|---|---|---|---|
| Eastward | $u_\theta = 0$<br>$u_\phi = \sin(\theta)$ | 0.0676 ±0.0141 | 1.60 ±0.041 | 0.710 ±0.015 |
| Westward | $u_\theta = 0$<br>$u_\phi = -\sin(\theta)$ | 0.0595 ±0.023 | 1.62 ±0.022 | 0.680 ±0.018 |
| Northward | $u_\theta = -\sin(\theta)$<br>$u_\phi = 0$ | 0.0817 ±0.020 | 1.54 ±0.011 | 0.680 ±0.0075 |
| Southward | $u_\theta = \sin(\theta)$<br>$u_\phi = 0$ | 0.0809 ±0.031 | 1.53 ±0.026 | 0.657 ±0.015 |
| North-East | $u_\theta = \sin(\theta)$<br>$u_\phi = \sin(\theta)$ | 0.0843 ±0.0022 | 3.22 ±0.031 | 1.38 ±0.019 |
| North-West | $u_\theta = \sin(\theta)$<br>$u_\phi = -\sin(\theta)$ | 0.0320 ±0.012 | 0.100 ±0.033 | 0.0417 ±0.014 |
| South-West | $u_\theta = -\sin(\theta)$<br>$u_\phi = -\sin(\theta)$ | 0.0980 ±0.023 | 3.18 ±0.048 | 1.37 ±0.018 |
| South-East | $u_\theta = -\sin(\theta)$<br>$u_\phi = \sin(\theta)$ | 0.0392 ±0.0083 | 0.120 ±0.027 | 0.043 ±0.017 |
| Rigid Body | See equation (4.14) | 0.205 ±0.028 | 0.832 ±0.039 | 0.811 ±0.019 |

Table 4.1: Synthetic flows used for testing. Similar results are found for flows constructed with $\cos\theta$.

## 4.3 Results

### 4.3.1 Atlantic

Having tested the methodology on synthetic examples, we now demonstrate an application of the methodology on the region under the equatorial Atlantic, using our optimal network size. We use CHAOS-8.1 SV from 1st January 2024, training models for five different seeds, and choose the one with the lowest RMSE as the 'preferred model'. Figure 4.7A shows $L_{SV}$, figure 4.7B shows $L_{FC}$, and figure 4.7C shows $L_{TOTAL}$, all plotted against iteration count, for one example seed. These loss curves show that the PINN has converged on a solution, and that the two loss components are of the same order of magnitude by the end of training, indicating that one loss term is not more important than the other. This vindicates our value of a weighting factor of $\lambda = 1000$. The $L_{FC}$ term has a much smaller amplitude in the first iteration than in any other iteration, which is not the typical shape of a loss curve, in which the loss at the first iteration is much larger than that of subsequent iterations. This is due to the magnitude of recovered flows in the first iteration being very close to zero, thereby satisfying the TG constraint but not the SV constraint, a behaviour that is penalised by the optimiser in subsequent iterations. While the loss terms are balanced by the end of the training run, the 'spikes' in the loss function indicate that the use of the weighting factor $\lambda$ may be too simple a method. Further work may be to implement adaptive weighting schemes, in which either the weighting factors are trainable parameters (McClenny and Braga-Neto, 2023) or the weighting factors are adjusted according to the model performance during training (Wang et al., 2021; Heydari et al., 2019).

The inferred flows are shown in figure 4.8A, the recovered SV in 4.8B, and the residuals shown in figure 4.8C. The mean absolute error (MAE) between the SV from the CHAOS-8.1 model and the SV output by the PINN is 97nT/year (Relative L1 Error of 2.57%), and the recovered flows show the expected westward drift present in all flow inversions to date, regardless of inversion methodology (Holme, 2007; Kloss and Finlay, 2019; Whaler et al., 2022).

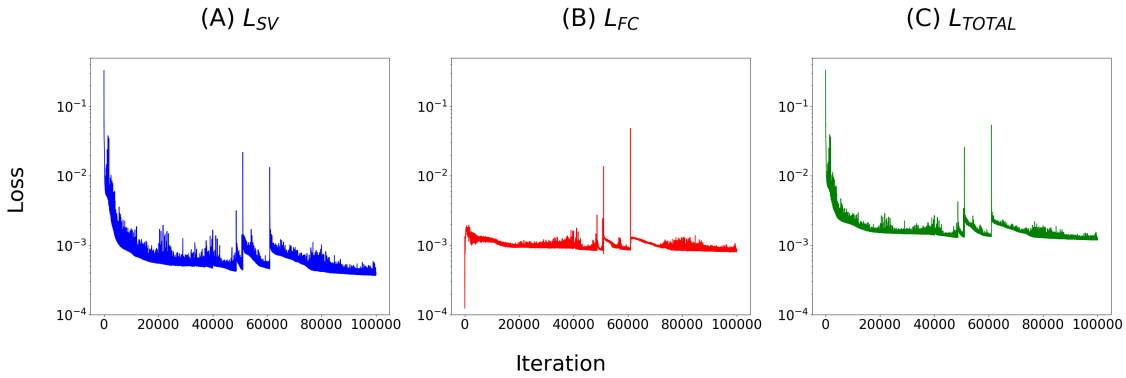



Figure 4.7: Loss plotted against iteration, (A) Data loss, $L_{SV}$, (B) TG Flow Constraint, $L_{FC}$, and (C) Total loss,$L_{TOTAL}$. The total loss does not decrease after 80,000 epochs, showing that the model has converged.

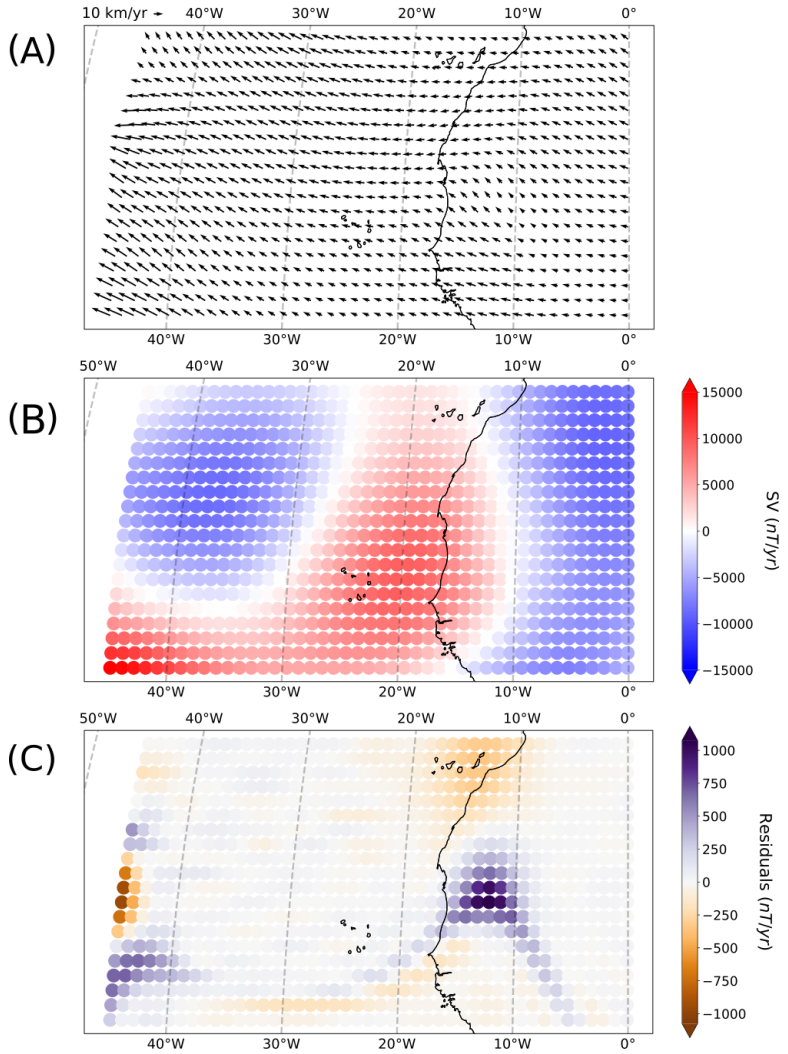



Figure 4.8: (A) Recovered Flow under the northern Atlantic, at 35°N to 5°N latitude, showing westward flow.
(B) Recovered SV
(C) Residuals between the SV from CHAOS-8.1 and the SV shown in (B).

### 4.3.2 Global Patchwork

In order to compare the flows recovered by the regional method with those recovered by global methods, a mosaic of regional flows is constructed from SV from 1st January 2024 using CHAOS-8.1, spanning a latitude space of within 1° of the geographic poles. We do not include the poles in our inversion, due to the coordinate singularity there. We also do not impose any constraints to impose boundary conditions in latitude or longitude. The surface of the core mantle boundary, at a radius of 3485km, is carved up into 56 overlapping latitude-longitude boxes spanning 30° by 55°, and 16 overlapping boxes spanning 25° by 55° for the boxes closest to the poles. Each of these boxes have grid points every 1° in both dimensions.

For each of the 72 boxes, a separate PINN is trained to find the flow that would generate the SV for that box only. The PINN for each box is trained for 100,000 iterations, and results obtained for 5 seeds. The model with the lowest $L_{SV}$ (SV RMSE) is taken as the 'preferred model' for that box. The 5° border is then cut off the results of each box, as in Figure 4.2, with the result being 56 20° by 45° boxes, and 16 15° by 45° boxes. The remainder of each box are stitched together, with no continuity conditions imposed at the edges of the boxes. The results go to within 6° of the pole. We did attempt to interpolate the flow in the overlapping edges to attempt to remove any flow discontinuities, but unfortunately this amplified those discontinuities, and so we have chosen to apply no additional smoothing.

The recovered flows, the recovered SV, and the residuals are shown in figure 4.9. The flows show the same large scale features as those found using global methodologies, demonstrating westward flow under the Atlantic, eastward flow under the Pacific, and the presence of the anticyclonic planetary gyre. This last feature is of particular note, as it demonstrates that this large scale feature is independent from the use of global methodologies, and is something that can be re-constructed from regional SV. The MAE between the SV from the CHAOS-8.1 model and the SV output by the PINN is 229nT/year (Relative L1 Error of 5.85%), with a standard deviation spatially of 251 nT/year and a mean TG residual of $7.7 \times 10^{-3}$ yr$^{-2}$. The residuals vary spatially, particularly in areas of increased SV intensity or complexity, and notably south of Africa where the maximum residual is 7578 nT/year. This high residual is consistent across seeds. Further investigation of this feature shows that the error is not consistent across time, peaking in 2024, and is not present when the TG assumption is not applied, perhaps indicating a failure of the TG assumption at this particular location. This is discussed in section 4.4.

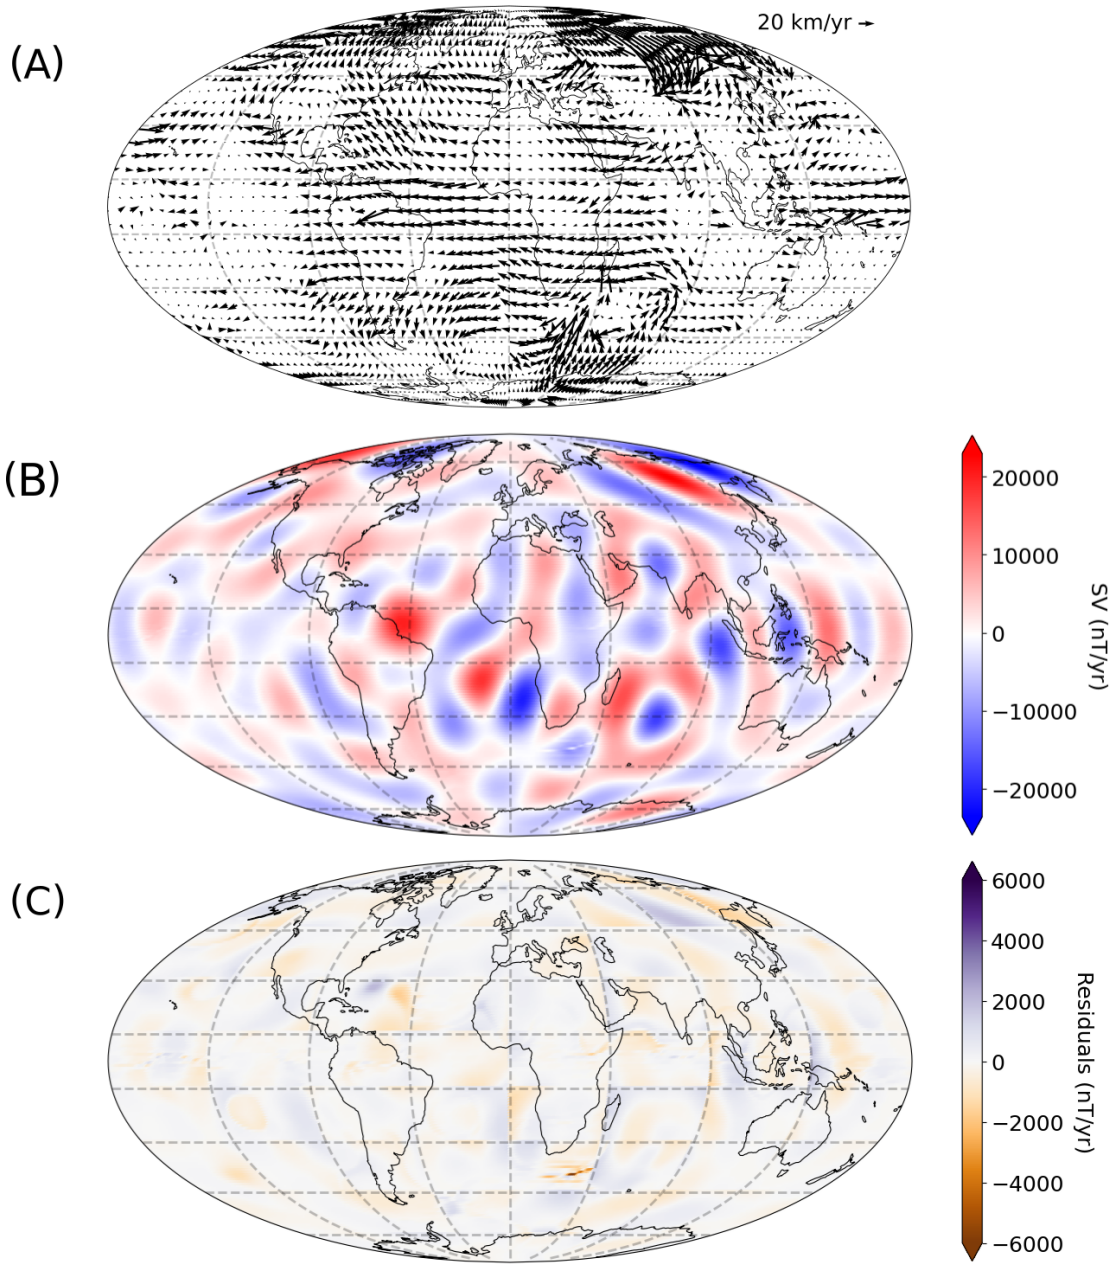



Figure 4.9: (A) Global patchwork of regional flows in January 2024. The thickness of the arrows indicate flow magnitude, and the arrows indicate direction.
(B) SV calculated from equation (4.1), using the recovered flows in the top panel.
(C) Residuals between the SV from CHAOS-8.1 and the SV shown in the middle panel.
The box boundaries (after the 5° border removal) are marked by grey dashed lines.

### 4.3.3 Equatorial Flows

To probe the broader flow properties of the equatorial region, the PINN methodology was applied to a 30° latitude by 360° longitude band, centred on the equator. SV values from CHAOS-8.1 were taken in one year increments from 2000.0 to 2024.0, and then five different models (one per seed) were trained for each of the 25 epochs, resulting in 125 models. For each time increment, the model from the five seeds which had the lowest RMSE between the input and output SV was chosen as the 'preferred model'. We use the same network size found in section 4.2.4 for this elongated box, as it was found empirically that larger network sizes did not give significantly lower values of SV RMSE, likely related to limitations of the TG assumption at low latitudes. Additionally, we do not impose periodicity between the longitudinal-edges of the elongated box.

The recovered azimuthal flow at the equator was extracted from the preferred models, and displayed in a time-longitude plot, shown in figure 4.10A. The area from 100° to 180° longitude shows a reversal in azimuthal flow direction, starting in 2010 and becoming more intense in 2017. This result was also found by Whaler et al. (2022), who associated the change in azimuthal flow direction with the 2017 geomagnetic jerk. By 2024.0, this eastward flow is still present. The area -180° to -70° longitude shows features of small-scale alternating flow direction in the area under central America, with four changes in flow direction seen in the area since 2000. This alternating pattern was previously found in Kloss and Finlay (2019) in their global flow inversion from ground- and satellite-based magnetic field measurements. Results from 2022 onwards show a recent change in the azimuthal flow direction in this area, from eastwards to westwards, which to our knowledge is a new result.

Figure 4.10B shows the recovered SV from the PINN, and figure 4.10C shows the residuals between the input and output SV. The residuals are generally small, with a MAE of 434 nT/year (Relative L1 Error of 9.98%), though the residuals vary slightly across longitude and time, with the largest residuals being at location of fast changes in azimuthal flow direction. Of particular note is a feature at -75° longitude, spanning from 2013 to 2018, where the maximum MAE is 6465 nT/year. This feature is not present when the TG assumption is not applied, once again indicating the possibility that this residual is due to a failure in the TG assumption, similar to that found in the global residuals shown in figure 4.9C. The residuals do not have any correlation to contours of $B_r/\cos\theta$, and most of the residuals appear to be due to a difference in intensity,

rather than structure.

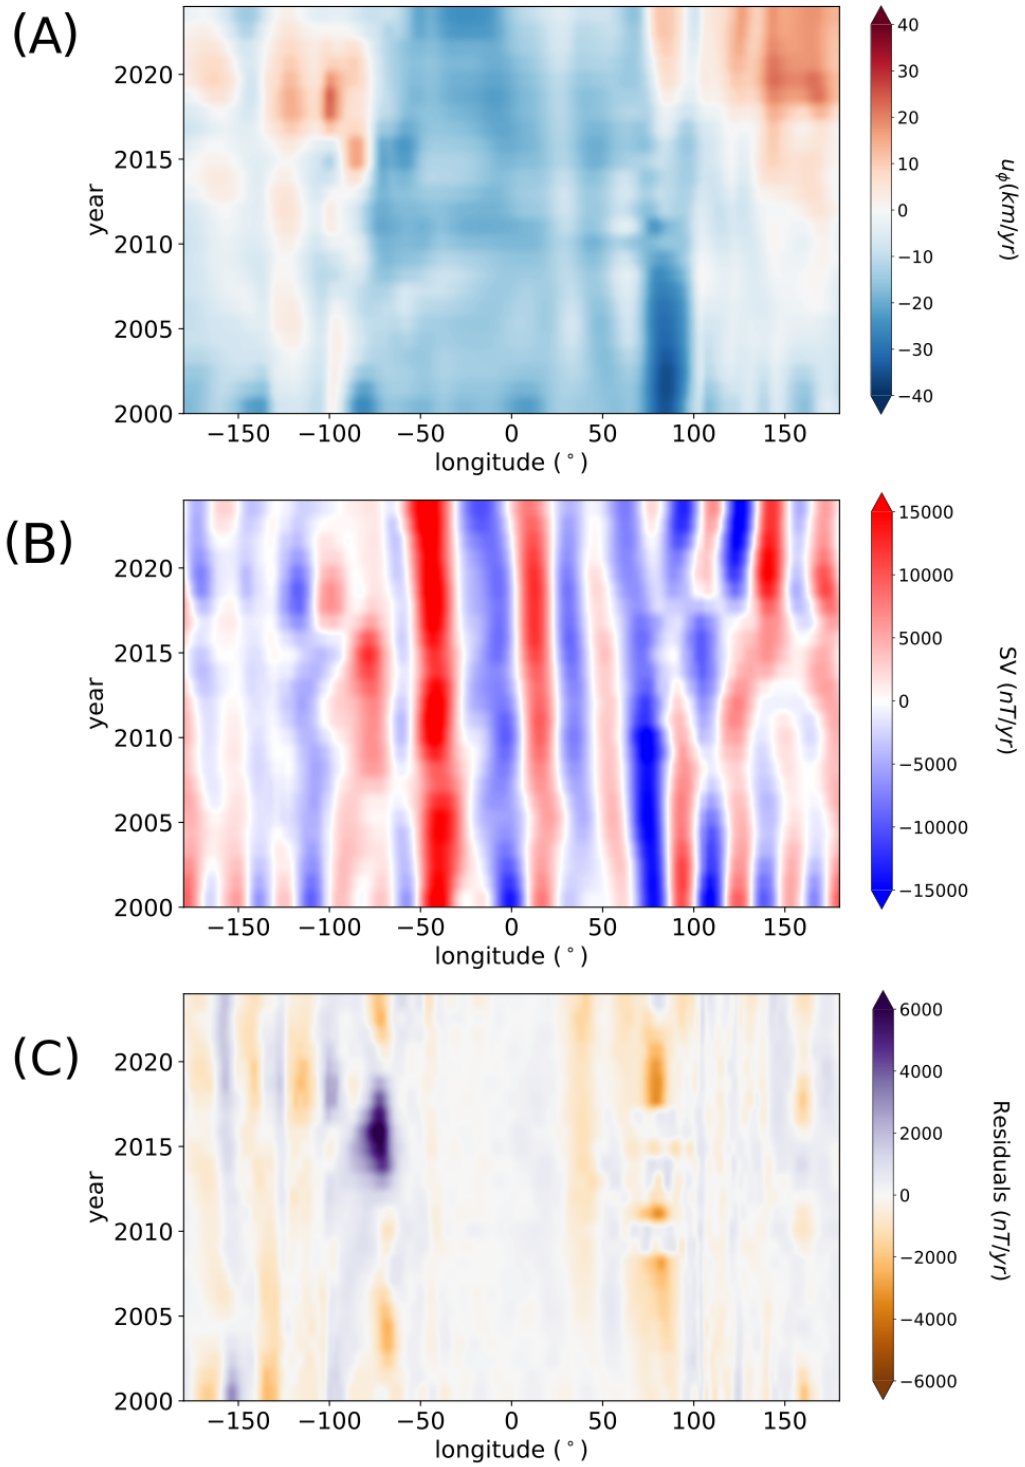



Figure 4.10: (A) Time-Longitude plot at the equator, spanning 2000.0 to 2024.0. Blue indicates westward flow, and red indicates eastward.
(B) SV calculated from equation (4.1), using the recovered flows in the top panel.
(C) Residuals between the SV from CHAOS-8.1 and the SV shown in the middle panel.

The two regions of azimuthal flow directional change indicated in figure 4.10 — Indonesia and Central America — were studied in more detail. Once again, for each time increment five models were trained, and the one with the lowest RMSE was chosen to be the 'preferred model'.

### Indonesia

The recovered flows under Indonesia in the time period of 2010.0-2024.0 are shown in figure 4.11, and the residuals for 2024.0 shown in figure 4.12. The flows show a predominantly westward, horizontally divergent flow structure beginning in 2010, increasing in magnitude until it becomes predominantly eastward flowing in 2020. Typical flow speeds are approximately 25 km/year, similar in magnitude to those found in Whaler et al. (2022). Results from 2015 onwards show an increased north-eastern flow magnitude at 150° to 165° longitude. The SV residuals show features that line up almost exactly with contours of $B_r/\cos\theta$, as shown in figure 4.12, the lines along which the flow is ambiguous. This is perhaps explained by considering equation (4.5), which shows that the induced SV is proportional to $B_r/\cos\theta$, and we might expect a correlation between the SV residuals and $B_r/\cos\theta$. The average MAE across the time period of 2010.0 - 2024.0 is 171 nT/year (Relative L1 Error of 2.65%), with a maximum of 255 nT/year in 2005 and a minimum of 120 nT/year in 2000.

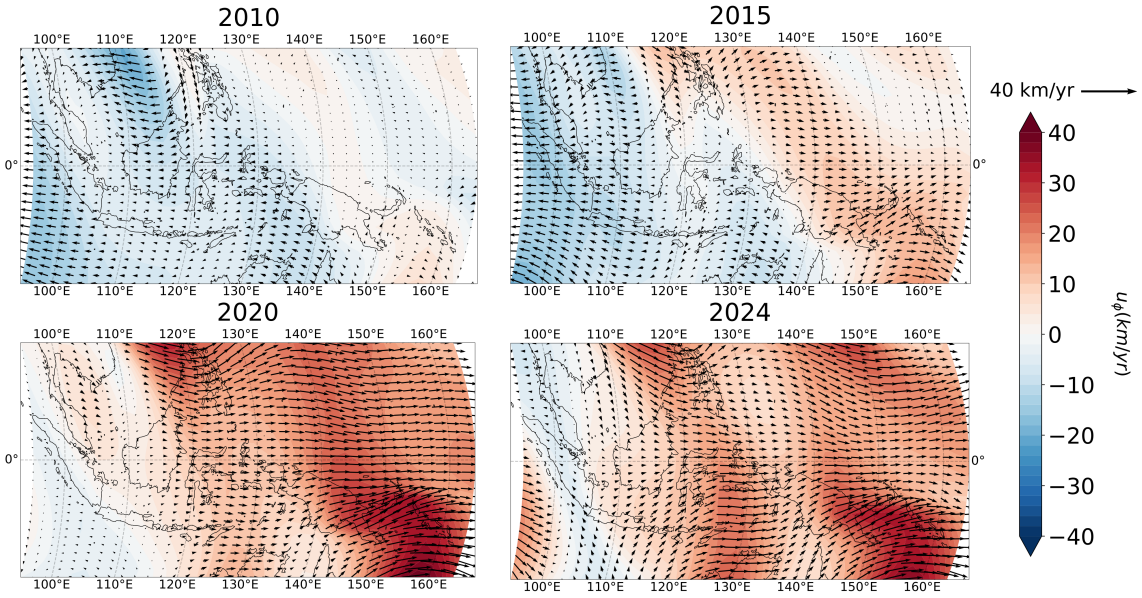



Figure 4.11: Changes in the horizontal flow under Indonesia during the period 2010.0 to 2024.0, at 10° above and below the equator.

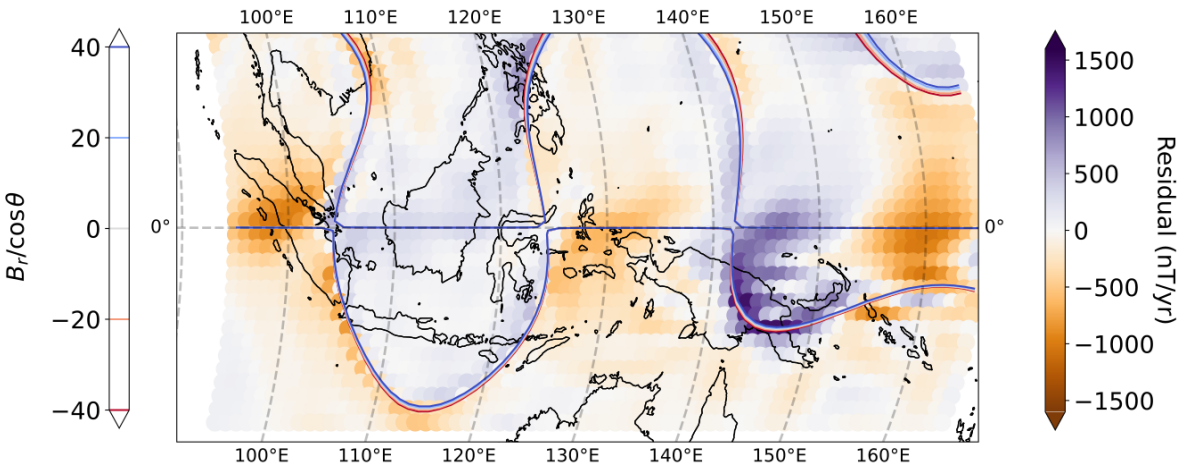



Figure 4.12: Residual between input SV and output SV for 2024.0, MAE of 157nT/year, with contours of $B_r/\cos\theta$ superimposed, truncated at $\pm$ 40 $nT$. Red lines indicate negative values of $B_r/\cos\theta$, and blue lines indicate positive values of $B_r/\cos\theta$. Latitude range is 10° above and below the equator.

### Eastern Pacific

Figure 4.13 shows three time increments from 2020.0 to 2024.0, which were chosen to investigate the recent change in flow direction in the Eastern Pacific, between longitudes of -150° to -75°. The recovered flows show a horizontally convergent flow structure in 2020, just under the coast of western South America, similar to those found in numerical simulations driven by heterogeneous CMB heat flow by Mound and Davies (2023), although the flow is strongly time-dependent and differs only a few years later. The intensity of the eastern portion of the recovered flow decreases from 2020, before becoming entirely westward by 2024. Once again, the residuals between input and output SV show features aligned with contours of $B_r/\cos\theta$. The average MAE across the time period of 2020.0 - 2024.0 is 201 nT/year (Relative L1 Error of 3.11%), with a maximum of 362 nT/year in 2020 and a minimum of 75 nT/year in 2024.

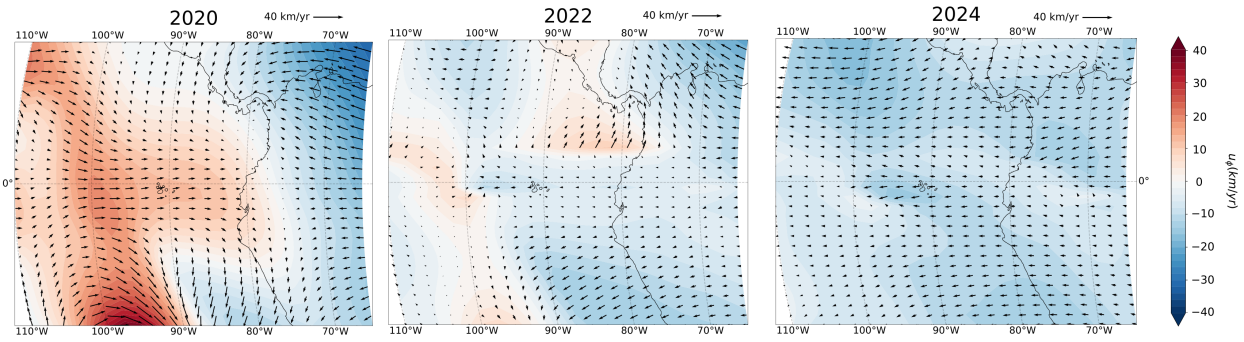



Figure 4.13: Changes under the Eastern Pacific during the period of 2020.0 to 2024.0, at 10° above and below the equator.

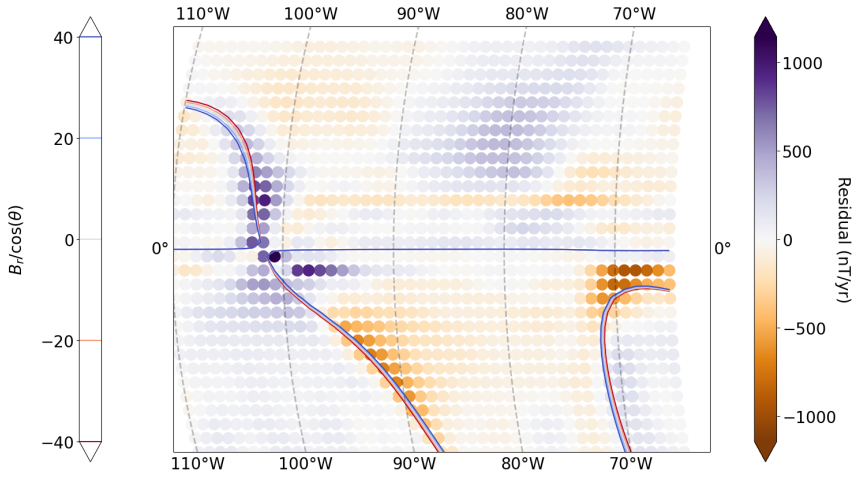



Figure 4.14: Residual between input SV and output SV for 2022.0, MAE of 201nT/year, with contours of $B_r/\cos\theta$ superimposed, truncated at $\pm$ 40 $nT$. Red lines indicate negative values of $B_r/\cos\theta$, and blue lines indicate positive values of $B_r/\cos\theta$. Latitude range is 10° above and below the equator.

### 4.3.4 High Latitude Jet near the North Pole

Global magnetic field models show patches of SV with alternating sign in the northern polar region, occurring close to the inner core tangent cylinder. Livermore et al. (2017) presented an explanation for this in the form of a strengthening high latitude jet of localised flow underneath Alaska and Siberia, which forms a part of the eccentric planetary gyre. The magnitude of this jet is not well constrained, as noted in Finlay et al. (2023); simpler, localised, models such as in Livermore et al. (2017) show a larger flow acceleration compared to those from more complex global models such as Gillet et al. (2022), which use priors to constrain non-uniqueness. It is worth noting that this choice of prior impacts the azimuthal flow speed around the tangent cylinder when using data assimilation methodologies (Rogers et al., 2025). To investigate where the flows from this regional method would fit into this picture, as well as to investigate the evolution of this jet over time, the PINN method was applied to a high-latitude band spanning from 60°N to 80°N latitude, and between 0° to 360° longitude. Once again 5 models were trained to find the 'preferred model', for increments every 5 years between 2000 and 2024, using the network size found in section 4.2.4.

The recovered flows in 2017 and 2024 are shown in figure 4.15. The PINN methodology recovers the high-latitude jet, increasing in azimuthal velocity beginning in 2005, before reaching a maximum velocity of 39 km/year in 2017 and decreasing in velocity afterwards. This change in intensity may be caused by a change in the discontinuous dynamics across the inner core tangent

cylinder, as suggested in Finlay et al. (2023), and could be associated with the evolution of the planetary gyre. Figure 4.16 shows the maximum azimuthal flow over time from this study (blue), the localised model of Livermore et al. (2017) (orange), and a global flow model that utilizes statistics from dynamo simulations as prior information (Gillet et al., 2022) (green). The PINN flow results show the larger acceleration favoured by the local flow model presented in Livermore et al. (2017), but with a decreased amplitude compared to the other local and global flow models. We also show that the jet has been decelerating since 2017, perhaps indicating that the jet has been a transient feature that is now subsiding. This result does not change with network size or choice of seed. The average MAE across the time period of 2000.0 - 2024.0 was 107 nT/year (Relative L1 Error of 1.79%), with a maximum of 173 nT/year in 2015 and a minimum of 49 nT/year in 2000.

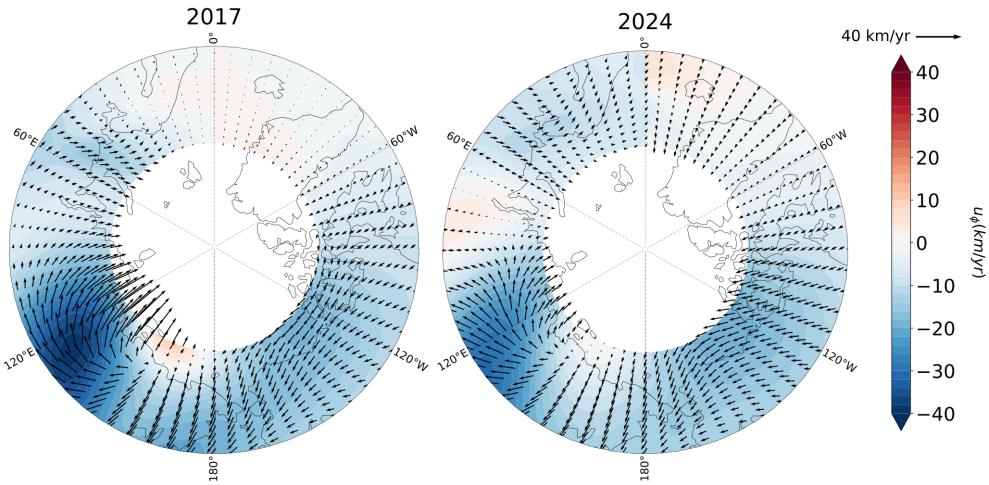



Figure 4.15: High Latitude jet in 2017.0 and 2024.0.

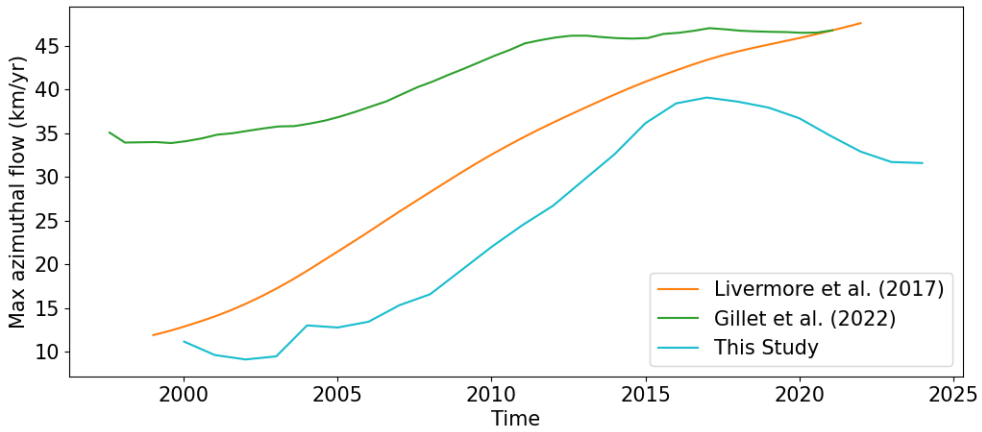



Figure 4.16: Maximum azimuthal flow over time, adapted from Finlay et al. (2023). Results from the model presented in Livermore et al. (2017) shown in orange, spanning a time period from 1999.0 to 2022.0, results from the model presented in Gillet et al. (2022) shown in green, spanning a time period from 1997.5 to 2021.0, and results from this study, spanning from 2000.0 to 2024.0, shown in blue.

## 4.4 Discussion

Our global mosaic model shown in figure 4.9 re-produces previously published large-scale flow structures, such as westward flow under the Atlantic, variable flow under Indonesia and Central America, and the large-scale gyre (Pais and Jault, 2008; Gillet et al., 2022; Kloss and Finlay, 2019). It is important to remember that our multiple regional flow models were all trained separately and had no knowledge of the SV outside the given region, but could still produce mostly continuous large-scale features across the box edges. The reproduction of the large-scale eccentric gyre indicates this is a data-driven feature, rather than it being an artefact of the global modelling approach. It is interesting to compare our results with the results shown in Buffett et al. (2016), where the authors found that high-latitude SV can be fit locally using waves, meaning there would be no need for a gyre to connect areas at high latitude to those at mid-latitude. Our models indicate that the high-latitude flow is connected to the mid-latitude flow, suggesting that the SV at high-latitudes is a flow feature, rather than a wave feature. This does not rule out high-latitude waves, but perhaps lends weight to the idea that there is a flow component due to the high-latitude SV.

The time-longitude plot at the equator, shown in figure 4.10, demonstrates azimuthal flow direction changes happening multiple times over a period of 24 years. There is good temporal continuation from one time-step to the next for the flow and recovered SV, despite there being no temporal regularisation applied to these PINN models. Instead, figure 4.10 shows a series of instantaneous snapshots over time, though there is temporal smoothing applied to the CHAOS-8.1 field model that provides the input SV. The regional flow inversions shown in figures 4.11 and 4.13 also highlight new features, particularly the recent change in flow direction under the eastern Pacific. This behaviour could be wave-driven, similar to the waves described in Gillet et al. (2022), but further work would be needed to confirm this. Recent changes are also shown in figure 4.16, in the decrease in azimuthal flow intensity at high latitudes in the Northern Hemisphere. Whether this change is due to a change in the dynamics at the inner core (Finlay et al., 2023), or due to other dynamics, is a possible area of further study.

While our study presents flows modelled with the tangentially geostrophic flow assumption, this framework makes changing the flow assumption, in principle, as simple as changing a single line of code. This allows for different flow assumptions to be tested easily and quickly. Additionally,

there is also no restriction on the shape of the region studied, as the loss function is a sum over a grid, which can be any shape. This provides the opportunity to study non-rectangular regions in their entirety, such as features of heterogeneity at the bottom of the mantle (e.g. Large Low-shear-Velocity Provinces, LLVPs (Panton et al., 2025)). However, care must be taken when choosing a flow assumption, as the residuals presented in this study suggest there are locations and times where the TG assumption is not valid. Additionally, due to the construction method of the CHAOS-8.1 model, our local values for the main field and SV at the CMB do still depend on the observational data everywhere. For a truly regional flow inversion, a regional geomagnetic model would need to be used, which is outside the scope of this study.

Regional core surface flow inversions have a wide range of potential applications, particularly in cases when the spatial distribution of data measurements is uneven, or when using paleomagnetic or historical data. The combination of this regional methodology with the SOLA (Subtractive Optimally Localised Averages, Hammer and Finlay (2018)) could be particularly productive, and may provide an opportunity for the use of historical/palaeomagnetic data for regional core flow inversions, as well as low inclination satellite orbits where data is restricted to mid- and low-latitudes, such as MSS-1 (Jiang et al., 2024).

## 4.5 Conclusion

This study presents a novel methodology for inverting regional flows at the core surface, using Physics Informed Neural Networks. We validate our method on synthetic flow examples, and then construct a global flow model from a mosaic of local flow models, re-producing the large-scale gyre and therefore indicating that this is a data-driven feature. We also investigate equatorial time dependence, and present regional flow models underneath the Atlantic, Indonesia, South America and at high latitudes in the Northern Hemisphere. We highlight new features, particularly the recent change in flow direction under the Eastern Pacific and the deceleration of the high-latitude jet from 2017.

## References

Alken, P., Thébault, E., Beggan, C.D., Amit, H., Aubert, J., Baerenzung, J., Bondar, T.N., Brown, W.J., Califf, S., Chambodut, A., Chulliat, A., Cox, G.A., Finlay, C.C., Fournier, A., Gillet, N., Grayver, A., Hammer, M.D., Holschneider, M., Huder, L., Hulot, G., Jager, T., Kloss, C., Korte, M., Kuang, W., Kuvshinov, A., Langlais, B., Léger, J.M., Lesur, V., Livermore, P.W., Lowes, F.J., Macmillan, S., Magnes, W., Mandea, M., Marsal, S., Matzka, J., Metman, M.C., Minami, T., Morschhauser, A., Mound, J.E., Nair, M., Nakano, S., Olsen, N., Pavón-Carrasco, F.J., Petrov, V.G., Ropp, G., Rother, M., Sabaka, T.J., Sanchez, S., Saturnino, D., Schnepf, N.R., Shen, X., Stolle, C., Tangborn, A., Tøffner-Clausen, L., Toh, H., Torta, J.M., Varner, J., Vervelidou, F., Vigneron, P., Wardinski, I., Wicht, J., Woods, A., Yang, Y., Zeren, Z., Zhou, B., 2021. International geomagnetic reference field: the thirteenth generation. Earth, Planets and Space 73, 49. doi:10.1186/s40623-020-01288-x.

Amit, H., 2014. Can downwelling at the top of the earth's core be detected in the geomagnetic secular variation? Physics of the Earth and Planetary Interiors 229, 110–121. URL: https://www.sciencedirect.com/science/article/pii/S0031920114000132, doi:https://doi.org/10.1016/j.pepi.2014.01.012.

Amit, H., Olson, P., 2004. Helical core flow from geomagnetic secular variation. Physics of the Earth and Planetary Interiors 147, 1–25. doi:10.1016/j.pepi.2004.02.006.

Amit, H., Pais, M.A., 2013. Differences between tangential geostrophy and columnar flow. Geophysical Journal International 194, 145–157. doi:10.1093/gji/ggt077.

Armstrong, J.A., Fletcher, L., 2019. Fast solar image classification using deep learning and its importance for automation in solar physics. Solar Physics 294, 80. doi:10.1007/s11207-019-1473-z.

Arzani, A., Wang, J.X., D'Souza, R.M., 2021. Uncovering near-wall blood flow from sparse data with physics-informed neural networks. Physics of Fluids 33, 071905. doi:10.1063/5.0055600. arXiv:2104.08249 [physics].

Backus, 1968. Kinematics of geomagnetic secular variation in a perfectly conducting core. Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences 263, 239–266.

Backus, Le Mouël, J., 1986. The region on the core—mantle boundary where a geostrophic velocity field can be determined from frozen-flux magnetic data. Geophysical Journal International 85, 617–628. doi:10.1111/j.1365-246X.1986.tb04536.x.

Bloxham, 1988. The dynamical regime of fluid flow at the core surface. Geophysical Research Letters 15, 585–588. doi:10.1029/GL015i006p00585.

Bloxham, Jackson, 1991. Fluid flow near the surface of earth's outer core. Reviews of Geophysics 29, 97–120. doi:10.1029/90RG02470.

Brown, W.J., Mound, J.E., Livermore, P.W., 2013. Jerks abound: An analysis of geomagnetic observatory data from 1957 to 2008. Physics of the Earth and Planetary Interiors 223, 62–76. URL: https://www.sciencedirect.com/science/article/pii/S0031920113000708, doi:10.1016/j.pepi.2013.06.001.

Buffett, B., Knezek, N., Holme, R., 2016. Evidence for MAC waves at the top of Earth's core and implications for variations in length of day. Geophysical Journal International 204, 1789–1800. URL: https://doi.org/10.1093/gji/ggv552, doi:10.1093/gji/ggv552, arXiv:https://academic.oup.com/gji/article-pdf/204/3/1789/1988933/ggv552.pdf.

Bullard, E.C., 1950. The transfer of heat from the core of the earth. Geophysical Journal International 6, 36–41. doi:10.1111/j.1365-246X.1950.tb02971.x.

Chen, Y., de Ridder, S.A.L., Rost, S., Guo, Z., Wu, X., Chen, Y., 2022. Eikonal tomography

with physics-informed neural networks: Rayleigh wave phase velocity in the northeastern margin of the tibetan plateau. Geophysical Research Letters 49, e2022GL099053. doi:10.1029/2022GL099053.

Constable, C., Constable, S., 2023. A grand spectrum of the geomagnetic field. Physics of the Earth and Planetary Interiors 344, 107090. URL: https://www.sciencedirect.com/science/article/pii/S0031920123001164, doi:https://doi.org/10.1016/j.pepi.2023.107090.

Finlay, Kloss, Olsen, N., Hammer, M.D., Tøffner-Clausen, L., Grayver, A., Kuvshinov, A., 2020. The chaos-7 geomagnetic field model and observed changes in the south atlantic anomaly 72, 156. doi:10.1186/s40623-020-01252-9.

Finlay, C.C., Gillet, N., Aubert, J., Livermore, P.W., Jault, D., 2023. Gyres, jets and waves in the earth's core. Nature Reviews Earth & Environment 4, 377–392. URL: https://www.nature.com/articles/s43017-023-00425-w, doi:10.1038/s43017-023-00425-w.

Friis-Christensen, E., Lühr, H., Hulot, G., 2006. Swarm: A constellation to study the earth's magnetic field. Earth, Planets and Space 58, 351–358. doi:10.1186/BF03351933.

Gillet, Gerick, F., Jault, D., Schwaiger, T., Aubert, J., Istas, M., 2022. Satellite magnetic data reveal interannual waves in earth's core. Proceedings of the National Academy of Sciences of the United States of America 119, e2115258119. doi:10.1073/pnas.2115258119.

Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks, in: Teh, Y.W., Titterington, M. (Eds.), Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, PMLR, Chia Laguna Resort, Sardinia, Italy. pp. 249–256. URL: https://proceedings.mlr.press/v9/glorot10a.html.

Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. MIT Press. http://www.deeplearningbook.org.

Gubbins, D., Hide, R., Osmaston, M.F., Runcorn, S.K., Creer, K.M., Jacobs, J.A., 1982. Finding core motions from magnetic observations. Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences 306, 247–254. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.1982.0084, doi:10.1098/rsta.1982.0084, arXiv:https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.1982.0084.

Hammer, M.D., Finlay, C.C., 2018. Local averages of the core–mantle boundary magnetic field from satellite observations. Geophysical Journal International 216, 1901–1918. URL: https://doi.org/10.1093/gji/ggy515, doi:10.1093/gji/ggy515, arXiv:https://academic.oup.com/gji/article-pdf/216/3/1901/27493010/ggy515.pdf.

Hansen, P.C., O'Leary, D.P., 1993. The use of the l-curve in the regularization of discrete ill-posed problems. SIAM Journal on Scientific Computing 14, 1487–1503. URL: https://doi.org/10.1137/0914086, doi:10.1137/0914086, arXiv:https://doi.org/10.1137/0914086.

Heydari, A.A., Thompson, C.A., Mehmood, A., 2019. Softadapt: Techniques for adaptive loss weighting of neural networks with multi-part loss functions. CoRR abs/1912.12355. URL: http://arxiv.org/abs/1912.12355, arXiv:1912.12355.

Holme, R., 2007. 8.04 - Large-Scale Flow in the Core. Elsevier, Amsterdam. p. 107–130. URL: https://www.sciencedirect.com/science/article/pii/B9780444527486001279, doi:10.1016/B978-044452748-6.00127-9.

Holme, R., 2015. 8.04 - large-scale flow in the core, in: Schubert, G. (Ed.), Treatise on Geophysics (Second Edition). second edition ed.. Elsevier, Oxford, pp. 91–113. URL: https://www.sciencedirect.com/science/article/pii/B978044453802400138X, doi:https://doi.org/10.1016/B978-0-444-53802-4.00138-X.

Jackson, A., Finlay, C., 2015. 5.05 - geomagnetic secular variation and its ap-

plications to the core, in: Schubert, G. (Ed.), Treatise on Geophysics (Second Edition). second edition ed.. Elsevier, Oxford, pp. 137–184. URL: https://www.sciencedirect.com/science/article/pii/B9780444538024000993, doi:https://doi.org/10.1016/B978-0-444-53802-4.00099-3.

Jiang, Y., Finlay, C.C., Olsen, N., Tøffner-Clausen, L., Yan, Q., Zhang, K., 2024. Macau scientific satellite-1 initial magnetic field model. Geophysical Research Letters 51, e2024GL112305. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2024GL112305, doi:https://doi.org/10.1029/2024GL112305, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2024GL112305. e2024GL112305 2024GL112305.

Jordan, Mitchell, 2015. Machine learning: Trends, perspectives, and prospects. Science 349, 255–260. doi:10.1126/science.aaa8415.

Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization URL: http://arxiv.org/abs/1412.6980, doi:10.48550/arXiv.1412.6980. arXiv:1412.6980 [cs].

Kloss, Finlay, 2019. Time-dependent low-latitude core flow and geomagnetic field acceleration pulses. Geophysical Journal International 217, 140–168. doi:10.1093/gji/ggy545.

Kloss, C., Finlay, C., Olsen, N., Tøffner-Clausen, L., Gillet, N., Grayver, A., 2024. The dtu geomagnetic field parent model and derived igrf-14 candidate models. URL: http://www.spacecenter.dk/files/magnetic-models/CHAOS-8/_downloads/1327fd0859c41d20bd663446966758ed/IGRF14_DTU_Technical_Note.pdf. technical Note for CHAOS-8.1, accessed 7 Jan 2025.

Langel, R.A., Estes, R.H., 1982. A geomagnetic field spectrum. Geophysical Research Letters 9, 250–253. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/GL009i004p00250, doi:https://doi.org/10.1029/GL009i004p00250, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/GL009i004p00250.

Le Mouël, J.L., 1984. Outer-core geostrophic flow and secular variation of earth's geomagnetic field. Nature 311, 734–735. URL: https://doi.org/10.1038/311734a0, doi:10.1038/311734a0.

Li, J., Lin, Y., Zhang, K., 2024. Dynamic mode decomposition of the core surface flow inverted from geomagnetic field models. Geophysical Research Letters 51, e2023GL106362. URL: https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2023GL106362, doi:10.1029/2023GL106362.

Livermore, P.W., Hollerbach, R., Finlay, C.C., 2017. An accelerating high-latitude jet in earth's core. Nature Geoscience 10, 62–68. doi:10.1038/ngeo2859.

Macmillan, S., 2007. Observatories, Overview. Springer Netherlands, Dordrecht. pp. 708–711. URL: https://doi.org/10.1007/978-1-4020-4423-6_224, doi:10.1007/978-1-4020-4423-6_224.

Mandea, Olsen, 2006. A new approach to directly determine the secular variation from magnetic satellite observations. Geophysical Research Letters 33. URL: https://onlinelibrary.wiley.com/doi/abs/10.1029/2006GL026616, doi:10.1029/2006GL026616.

McClenny, L.D., Braga-Neto, U.M., 2023. Self-adaptive physics-informed neural networks. Journal of Computational Physics 474, 111722. URL: https://www.sciencedirect.com/science/article/pii/S0021999122007859, doi:https://doi.org/10.1016/j.jcp.2022.111722.

Mound, J.E., Davies, 2023. Longitudinal structure of earth's magnetic field controlled by lower mantle heat flow. Nature Geoscience 16, 380–385. doi:10.1038/s41561-023-01148-9.

Olsen, N., Floberghagen, R., 2018. Exploring geospace from space: the swarm satellite constellation mission. Space Research Today 203, 61–71. URL: https://www.sciencedirect.com/

science/article/pii/S1752929818300884, doi:https://doi.org/10.1016/j.srt.2018.11.017.

Pais, M.A., Jault, D., 2008. Quasi-geostrophic flows responsible for the secular variation of the earth's magnetic field. Geophysical Journal International 173, 421–443. URL: https://doi.org/10.1111/j.1365-246X.2008.03741.x, doi:10.1111/j.1365-246X.2008.03741.x, arXiv:https://academic.oup.com/gji/article-pdf/173/2/421/6078007/173-2-421.pdf.

Panton, J., Davies, J.H., Koelemeijer, P., Myhill, R., Ritsema, J., 2025. Unique composition and evolutionary histories of large low velocity provinces. Scientific Reports 15, 4466. URL: https://doi.org/10.1038/s41598-025-88931-3, doi:10.1038/s41598-025-88931-3.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. Pytorch: An imperative style, high-performance deep learning library. URL: https://arxiv.org/abs/1912.01703, arXiv:1912.01703.

Raissi, Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378, 686–707. doi:10.1016/j.jcp.2018.10.045.

Raissi, Yazdani, A., Karniadakis, G.E., 2020. Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. Science (New York, N.Y.) 367, 1026–1030. doi:10.1126/science.aaw4741.

Rankine, W.J.M., 1872. A manual of applied mechanics. Charles Griffin and Company.

Rau, S., Christensen, U., Jackson, A., Wicht, J., 2000. Core flow inversion tested with numerical dynamo models. Geophysical Journal International 141, 485–497. URL: https://doi.org/10.1046/j.1365-246x.2000.00097.x, doi:10.1046/j.1365-246x.2000.00097.x, arXiv:https://academic.oup.com/gji/article-pdf/141/2/485/1826494/141-2-485.pdf.

Roberts, P.H., Scott, S., 1965. On analysis of the secular variation. Journal of geomagnetism and geoelectricity 17, 137–151. doi:10.5636/jgg.17.137.

Rogers, H.F., 2022. Improving inversions of regional outer core surface flow models through the application of spherical slepian function URL: https://era.ed.ac.uk/handle/1842/39295, doi:10.7488/era/2546. accepted: 2022-08-02T10:34:07Z.

Rogers, H.F., Gillet, N., Aubert, J., Personnettaz, P., Mandea, M., 2025. Effects of geodynamo priors and geomagnetic data on inverted core surface flows. Physics of the Earth and Planetary Interiors , 107323URL: https://www.sciencedirect.com/science/article/pii/S0031920125000172, doi:https://doi.org/10.1016/j.pepi.2025.107323.

Schiassi, E., De Florio, M., Ganapol, B.D., Picca, P., Furfaro, R., 2022. Physics-informed neural networks for the point kinetics equations for nuclear reactor dynamics. Annals of Nuclear Energy 167, 108833. doi:10.1016/j.anucene.2021.108833.

Schwaiger, T., Jault, D., Gillet, N., Schaeffer, N., Mandea, M., 2023. Local estimation of quasi-geostrophic flows in earth's core. Geophysical Journal International 234, 494–511. URL: https://doi.org/10.1093/gji/ggad089, doi:10.1093/gji/ggad089, arXiv:https://academic.oup.com/gji/article-pdf/234/1/494/50509990/ggad089.pdf.

Shakespeare-Rees, N., Livermore, P.W., Davies, C.J., Rogers, H.F., Brown, W.J., Beggan, C.D., Finlay, C.C., 2025. Local flow estimation at the top of the earth's core using physics informed neural networks. Physics of the Earth and Planetary Interiors 367, 107424. URL: https://www.sciencedirect.com/science/article/pii/S0031920125001189, doi:https://doi.org/10.1016/j.pepi.2025.107424.

Sola, J., Sevilla, J., 1997. Importance of input data normalization for the application of neu-

ral networks to complex industrial problems. IEEE Transactions on Nuclear Science 44, 1464–1468. doi:10.1109/23.589532.

Wang, S., Teng, Y., Perdikaris, P., 2021. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. SIAM Journal on Scientific Computing 43, A3055–A3081. URL: https://doi.org/10.1137/20M1318043, doi:10.1137/20M1318043, arXiv:https://doi.org/10.1137/20M1318043.

Whaler, K.A., 1980. Does the whole of the earth's core convect? Nature 287, 528–530. doi:10.1038/287528a0.

Whaler, K.A., Hammer, M.D., Finlay, C.C., Olsen, N., 2022. Core surface flow changes associated with the 2017 pacific geomagnetic jerk. Geophysical Research Letters 49, e2022GL098616. doi:10.1029/2022GL098616.

Zhang, K., 2023. A novel geomagnetic satellite constellation: Science and applications. Earth and Planetary Physics 7, 4–21. URL: https://www.eppcgs.org/en/article/doi/10.26464/epp2023019, doi:10.26464/epp2023019.

# Chapter 5

# Forecasting Secular Variation using Physics-Informed Neural Networks

Naomi Shakespeare-Rees, Philip W. Livermore, Christopher J. Davies, Hannah F. Rogers, William J. Brown, Ciarán D. Beggan, Christopher C. Finlay

Published in *Earth, Planets and Space* as Shakespeare-Rees et al. (2026). Minor editorial changes have been made for consistency with the thesis.

## Abstract

In response to the call for candidate models for the 14th generation of the International Geomagnetic Reference Field (IGRF) by the Geomagnetic Field Modeling Working Group (V-MOD) of the International Association of Geomagnetism and Aeronomy (IAGA), we present the University of Leeds candidate model for the forecast of the average Secular Variation (SV) for 2025–2030. Our approach consists of inverting the geomagnetic field model CHAOS-7.18 using Physics-Informed Neural Networks to produce two global mesh-free models from (i) a mosaic of independent regional flows and (ii) a single global flow representation. The magnetic field is then advected under the assumption of steady core flow over a five-year period, and the average SV over 5 years is taken to construct the forecast. We validate our approach using hindcasts for the IGRF-13 time period (2020-2025) and benchmark our methodology against the inferred SV from CHAOS-7.18. Our field models constructed from regional flows show reduced RMS misfit relative to the field from the CHAOS-7.18 model at each yearly timestep, compared to the other candidate models from IGRF-13 in the hindcast, both at the Core Mantle Boundary and at the Earth's Surface. We then present our IGRF-14 candidate forecast for the period 2025-2030, derived from the regional method, and discuss possible improvements to this method for future IGRF submissions.

## 5.1 Introduction

The Earth's main magnetic field, powered by fluid motions in the Earth's outer core that drive a self-sustaining dynamo (Bullard, 1950), changes over a range of timescales. The observed timescales span years to millennia (Constable and Constable, 2023) and temporal variations in the geomagnetic field on timescales of years to millennia are termed Secular Variation (SV) (Jackson and Finlay, 2015). The geomagnetic field has been measured by networks of ground-based observatories since 1837 (Macmillan, 2007), and has been further supported by continuous satellite measurements since 1999 (Friis-Christensen et al. (2006), Olsen and Kotsiaros (2011), Olsen and Floberghagen (2018), Olsen et al. (2020), Zhang (2023), Jiang et al. (2024)). Temporal variations in the geomagnetic field are significant for both industry and scientific research across various fields, including navigation, satellite operations, and space weather hazards.

To try to understand and predict changes in the geomagnetic field, models, such as the International Geomagnetic Reference Field (IGRF), are created. Commonly, these models are constructed in spectral space with global basis functions, such as spherical harmonics, constrained by measurements from satellites, observatories and other surveys. A benefit to using spherical harmonic global basis functions is the ability to analytically downward continue the field model to the Core-Mantle Boundary (CMB). The IGRF is a collaborative effort overseen by the Division V Geomagnetic Field Modelling Working Group (V-MOD) of the International Association of Geomagnetism and Aeronomy (IAGA), which is updated every five years. Intervening epochs are forecast through an average SV model, itself derived from multiple candidate models. The previous edition of IGRF is the 13th iteration (Alken et al., 2021b), which contains a definitive main field model for 2015.0, a main field model for epoch 2020.0, and a predicted secular variation for 2020.0 to 2025.0.

The construction of the IGRF has driven advances in geomagnetic field modelling, evidenced by the variety of methods used to model the field. Responding to previous IGRF calls, candidate methodologies for secular variation have included constructing models from ground and satellite measurements (Alken et al. (2021a), Finlay et al. (2020), Rother et al. (2021), Huder et al. (2020)), data assimilation techniques (Fournier et al. (2021), Minami et al. (2020), Sanchez et al. (2020)), Kalman-filter based models (Tangborn et al. (2021), Baerenzung et al. (2020)), interpolation (Petrov and Bondar, 2021), diffusion based models (Metman et al., 2020a), boot-

strapping (Pavón-Carrasco et al., 2020) and spectral analysis (Wardinski et al., 2020). All of these models, except for the diffusion based model presented by Metman et al. (2020b), assume that the observed SV is produced by advection of the fluid in the outer core.

A frequently used assumption when inverting the geomagnetic field to map flows at the edge of the outer core is to assume that on sufficiently short timescales, magnetic diffusion can be considered negligible - this is the *frozen flux* assumption presented in Roberts and Scott (1965). Under this simplification, the magnetic field lines can be thought to be 'frozen' within moving liquid packets at the surface of the outer core, effectively making the change in the magnetic field a tracer of the flow (Bloxham and Jackson, 1991). Non-uniqueness is further minimised by assuming large scale flow, wherein small scale flow structures are penalised (Bloxham, 1988). Flow features inferred from previous studies include an eccentric, anticyclonic, planetary gyre (Pais and Jault, 2008), changes in the azimuthal flow direction in areas in the equatorial region (Whaler et al., 2022; Li et al., 2024), and a localised jet under the Bering Strait (Livermore et al., 2017).

For our IGRF-14 SV forecast, we employ a steady advection-only core flow approach. A candidate model for IGRF-13 based on the steady-flow advection of the outer core flow was presented by Brown et al. (2021), with two candidate models of IGRF-12 also being based on steady core flow advection (Hamilton et al., 2015). However, in contrast to prior approaches, constructed using global basis functions, we implement two novel aspects based on neural networks: (i) a global model created by a mosaic of regional flows individually and independently inferred by a fit to a local realisation of a global geomagnetic field model, and (ii) a global model, inferred through a single global fit to a field model. Our regional approach provides an independent method to check global flow analyses, providing complementary information for the global candidate models, as a regularised large-scale global inversion may act to smooth out important local flow features. Additionally, in a model dependent on global basis functions, any point on the CMB will be influenced non-locally by the SV at all the other locations on the CMB, but in a regional flow inversion the same point only depends on the other points in that local region. This is important because, although the quality and coverage of geomagnetic measurements from satellites and ground-based observatories are generally good, satellite data are more susceptible to contamination from external fields in auroral regions, while the spatial distribution of the observatory network remains highly uneven. Our regional methodology can probe the

reliability of local features in these areas; we can also be sure that flow maps constructed in non-polar regions are not contaminated by polar data with significant unmodelled non-internal signal.

Local core flow inversions have been attempted in previous studies. Rogers (2022) studied regional variations using spherical Slepian functions. This technique produced better separation of SV at the Earth's surface compared to spherical harmonics but reliable local flow separations and SV separations at the CMB were not achieved. Schwaiger et al. (2023) presented a local core flow inversion methodology based on pointwise inversion, which was able to reproduce the main features found in global core flow studies, but it was found that additional smoothing was required to prevent unreliable re-construction of small-scale flows. They also found that their results heavily depended on what prior they used to reduce the non-uniqueness. Our approach infers regional flows through the use of Physics-Informed Neural Networks (PINNs), a class of machine learning models, described in Shakespeare-Rees et al. (2025).

Machine learning refers to a collection of statistical techniques used to leverage data in order to complete a task, allowing the user to extract knowledge and draw inferences (Jordan and Mitchell, 2015). Neural Networks (NN), a machine learning architecture, consist of layers of artificial neurons that can process information and learn complex representations of information. PINNs, first proposed in Raissi et al. (2019), are a class of neural network that integrate mathematical descriptions of physical laws as soft constraints in order to solve forward and inverse problems. PINNs have been used in a diverse range of fields such as fluid mechanics (Raissi et al., 2020), medicine (Arzani et al., 2021), nuclear physics (Schiassi et al., 2022) and seismology (Chen et al., 2022). PINNs have been recently used in core flow and length of day (LOD) analysis by Li et al. (2024), albeit only in a global flow inversion framework.

The rest of this paper is organised as follows: in Section 5.2 we describe our core flow inversion methodology using PINNs, and then detail our advection scheme which produces the SV forecast in Section 5.3. In Section 5.4, we show results for hindcasting, benchmarking our hindcast models against IGRF-13 candidate models, and then present our candidate SV forecast for IGRF-14 in Section 5.5. Finally, we discuss our results and areas of improvement for our methodology in Section 5.6.

## 5.2 Core Flow Inversion using PINNs

Modelling the Earth's core as a sphere described using spherical coordinates $(r, \theta, \phi)$, we assume that the flow, $\boldsymbol{u}$, within the sphere obeys a non-penetration boundary condition at the edge of the core with vanishing radial flow, $u_r = 0$, at the CMB. The radial magnetic field, $B_r$, can then be related to the horizontal fluid motion at the edge of the core via the radial component of the induction equation, which assuming frozen flux, can be written as:

$$\frac{\partial B_r}{\partial t} = -\nabla_H \cdot (\boldsymbol{u} B_r), \tag{5.1}$$

in which $\boldsymbol{u} = [0, u_\theta, u_\phi]$ is the flow that is sought in the inversion methodology, and $\nabla_H$ is the horizontal gradient operator.

In spherical coordinates, the total flow $\boldsymbol{u}$ can be decomposed into toroidal and poloidal flows

$$\boldsymbol{u} = \nabla \times T(\theta, \phi)\mathbf{r} + \nabla_H(rP(\theta, \phi)), \tag{5.2}$$

whose horizontal components can be written (Holme, 2015),

$$\boldsymbol{u}_T = \left( \frac{1}{\sin\theta} \frac{\partial T}{\partial \phi}, -\frac{\partial T}{\partial \theta} \right), \tag{5.3}$$

$$\boldsymbol{u}_P = \left( \frac{\partial P}{\partial \theta}, \frac{1}{\sin\theta} \frac{\partial P}{\partial \phi} \right). \tag{5.4}$$

Therefore, we seek toroidal ($T$) and poloidal ($P$) scalar functions which define the flow in our flow inversion. Our inversion methodology consists of two Fully Connected Neural Networks (FC-NNs) working in parallel: one to describe the toroidal scalar function $T$ and the other to describe the poloidal scalar function $P$. A schematic for this is shown in Figure 5.1. FC-NNs are composed of layers of nodes, including an input layer, one or more hidden layers, and an output layer (Jordan and Mitchell, 2015). Each node is interconnected with others and is associated with a specific weight and bias. The output of a node is calculated by applying a non-linear function—known as the activation function—to the weighted input. During a process known as training, the network adjusts its weights and biases to improve performance (Goodfellow et al., 2016). Once the network weights are determined, the flows are then described using Equations (5.3) and (5.4).

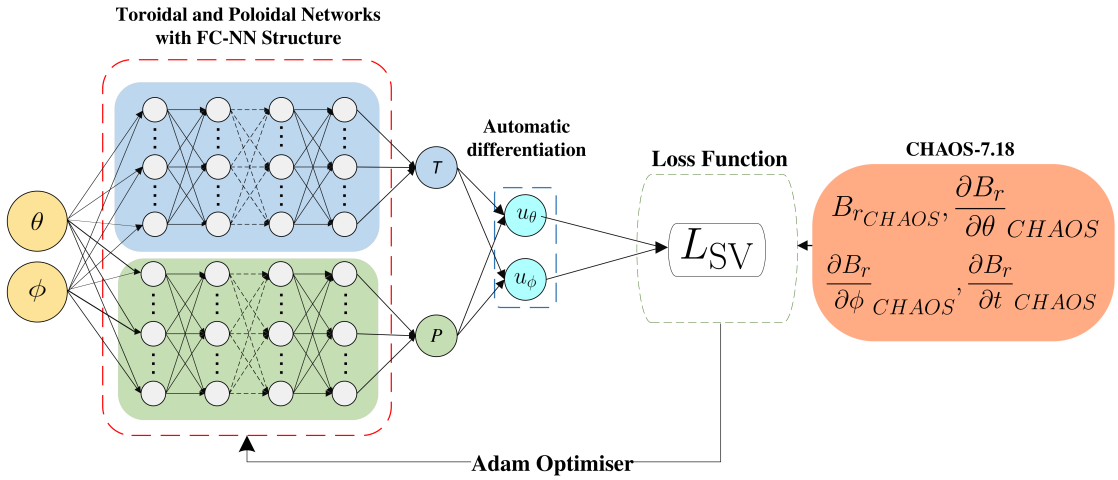



Figure 5.1: Schematic of the PINN used in this study, where $\theta, \phi$ are co-latitude and longitude, $T$ and $P$ are the toroidal and poloidal scalars, and $u_\theta$ and $u_\phi$ are the components of the horizontal flow. Yellow indicates inputs to the network, the two parallel FC-NNs are blue and green, loss function in white, data inputs to the loss function in orange, and the outputs are cyan.

The networks are trained using the optimisation method Adam (Kingma and Ba, 2017), which updates both networks concurrently. The PINN receives a grid of points defined by $(\theta, \phi)$ at a radius of 3485 km. It then analytically computes the toroidal and poloidal scalar value at each of those points, then calculates the spatial derivatives of these scalars using Pytorch's built-in differentiation engine (Paszke et al., 2017). These derivatives are then used in Equations 5.3 and 5.4 to produce $u_\theta, u_\phi$. The loss term in the PINNs used for this submission consists of the MSE (Mean Squared Error) between the SV recreated by the PINN-based flows and the SV from the CHAOS-7.18 model:

$$L_{\mathrm{SV}} = \frac{1}{N}\sum_{i}^{N}\left[\frac{\partial B_r}{\partial t}_{PINN} - \frac{\partial B_r}{\partial t}_{CHAOS}\right]^2\Bigg|_{\theta_i,\phi_i} \tag{5.5}$$

where $N$ is the number of points, and the SV calculated from the horizontal flows is

$$\frac{\partial B_r}{\partial t}_{PINN} = -\frac{1}{r}\left(\frac{u_\phi}{\sin\theta}\frac{\partial B_r}{\partial \phi}_{CHAOS} + u_\theta \frac{\partial B_r}{\partial \theta}_{CHAOS}\right) - B_{r_{CHAOS}}(\nabla_H \cdot \boldsymbol{u}). \tag{5.6}$$

The Adam optimizer back-propagates the loss through the network, updating weights and biases to minimize the loss via gradient descent over 5,000 iterations. However, empirical results show that $L_{SV}$ as a function of the model parameters exhibits complex behaviour, leading the optimiser to only converge to local minima (Shakespeare-Rees et al., 2025). We

therefore choose multiple different seeds – equivalent to varying initial weights – for independent training, from which we select the model that achieves the lowest loss. Given that our dataset is of order 1000-10000 points – modest by machine learning standards – we are able to use all the data in each iteration of training. The PINN is constructed using SciANN (Haghighat and Juanes, 2021), a TensorFlow/Keras wrapper for physics-informed machine learning, with straightforward implementation. We highlight that other approaches to constructing PINNs are available, such as the PyTorch approach described in Shakespeare-Rees et al. (2025).

In order to optimise training, a number of hyper-parameters need to be chosen. The learning rate – governing how rapidly the weights change with each iteration – was chosen to be the standard value of $10^{-3}$, as it was found empirically that this resulted in the lowest SV misfit. Our choice of activation function was swish (Ramachandran et al., 2017), due to its better performance when compared to other activation functions such as Tanh and ReLU. The size of the networks – the number of hidden layers and the number of neurons per layer – is an important parameter to choose. If the network is too small, then it will have a very limited functional representation, resulting in a flow too smooth to fit the SV data. If it is too large, then it could result in flows that fit the SV but have spurious small-scale features. We therefore choose the optimal size of network through sampling in model space, finding the size that produces the minimum SV residual but has no spurious small-scale flows, through the trade-off method detailed in Shakespeare-Rees et al. (2025), which we briefly describe here. We determine the optimal network size by balancing the quality of data fit against the complexity of the resulting flow, without applying explicit regularisation. To identify suitable configurations, we evaluate networks under two regions: the Atlantic, where simple westward flow is expected (Holme, 2007), and the South China Sea, which exhibits more complex, divergent flow behaviour (Whaler et al., 2022). For each network architecture, we compute the SV root mean squared error (RMSE) and the complexity as measured by the average of the squared second spatial derivative of the flow (Bloxham, 1988), evaluated on the SV data grid, repeating the process across five random seeds to assess variability and robustness. These results are then visualised as trade-off curves, from which we identify the 'knee' point—representing the best compromise between accuracy and simplicity. Due to the relatively small size of the dataset, cross-validation was not employed, as it would require partitioning the data during training and risk degrading model performance. Using this method, we found an optimal size of 10 layers of 150 neurons. Although the PINN methodology allows us to apply additional physics-based flow constraints, we chose not to do

this and to present the simplest possible model approach for this submission. Additionally, while we do not explicitly apply smoothing to the flows, there is smoothing inherent to the modelling approach due to representing the flow using a finite network as specified above.

### 5.2.1 Training Data

The training procedure for the PINN methodology requires values for the SV, as well as values for the horizontal spatial derivatives of the main field. We use the core field component of the geomagnetic field model CHAOS-7.18, an extension of the CHAOS-7 model presented in Finlay et al. (2020), projected onto a grid in spherical coordinates. CHAOS-7.18 spans from January 1999 to June 2024, and is built from a combination of satellite and ground-based observatory data. This was deemed a suitable choice of model, as it has been used in multiple core flow inversion studies (eg. Gillet et al. (2022)). CHAOS-7.18 is temporally regularised to reduce non-uniqueness by penalising the second time derivative at the endpoints, and the third time derivative throughout. The time-dependent internal field model is defined up to spherical harmonic degree 20, but only the coefficients for the main field and SV up to degree 8 are used for the SV forecast, this being the degree to which the SV forecast is required in IGRF. Truncation of the secular variation at higher spherical harmonic degrees was explored, but it did not lead to improved performance in hindcasts.

We tested two approaches to core flow inversion: a 'Regional' method and a 'Global' method. For the regional approach, a mosaic of regional flow models are constructed from SV from 1st January 2020 using CHAOS-7.18. The surface of the core mantle boundary, at a radius of 3485 km, is carved up into 56 overlapping latitude-longitude boxes each spanning 30° by 55°, and 16 overlapping boxes spanning 25° by 55° for the boxes closest to the poles. The boxes extend to within 5° of the poles. The edges of these boxes are shown by the grey lines in Figure 5.2. Each of these boxes have grid points every 1° in both dimensions. We note that alternative partitions of the CMB, such as overlapping circular regions or tessellated triangles, are also possible, but these were considered to be beyond the scope of the present study. For each of the 72 boxes, a separate PINN is trained to find the flow that would generate the SV for that box only. The PINN for each box is trained for 5000 iterations, and results obtained for 5 seeds. The model with the lowest $L_{SV}$ (SV MSE) is taken as the 'preferred model' for that box. A 5° border is then cut off the results of each box, as in Figure 5.3, with the result being 56 boxes of 20° by 45°, and 16 boxes of 15° by 45°. Sensitivity tests during development indicated that a 5° border was

sufficient to maintain continuity between adjacent regions without introducing notable artefacts. Increasing the border produced only marginal improvements in transition smoothness, with no significant changes in overall flow patterns or fine-scale features, even when the border size was doubled. The remainder of each box are stitched together, with no continuity conditions imposed at the edges of the boxes. For use in the prediction methodology, the $\boldsymbol{u}$ and $\nabla \cdot \boldsymbol{u}$ values were interpolated from the training grid, using cubic spline interpolation onto the same Gauss-Legendre-Fourier grid as the advection method in Section 5.3, as this gave better hindcast results than simply evaluating the model directly onto the Gauss-Legendre-Fourier grid.

For the global flow inversion model, we projected the field coefficients of CHAOS-7.18 on one large latitude-longitude box covering most of the CMB, aiming to compare a model incorporating all available spatial data, with our regional model incorporating the same data in regional boxes. The box extends to within 5° of the poles, and from -180° to 180° in longitude. We do not assume any regularity conditions for this approach, and we do not impose periodicity between the longitudinal-edges of the box. CHAOS-7.18 is projected onto the same Gauss-Legendre grid as the regional inversion. We did not add a border to the edges for the global flow inversion, as they would have extended over the poles.

To optimise model performance and reduce training time, the CHAOS-7.18 input and outputs of the network are re-scaled so that their magnitudes are all about 1 (Sola and Sevilla, 1997). This is achieved by expressing $B_r$ in $\mu$T, time in 0.1 yr, and length in km. Typical values of $B_r$ at the CMB are $500\mu$T and so typical values of SV are then 1 $\mu$T/0.1 yr in these units and flow speeds are typically 1 km/0.1 yr.

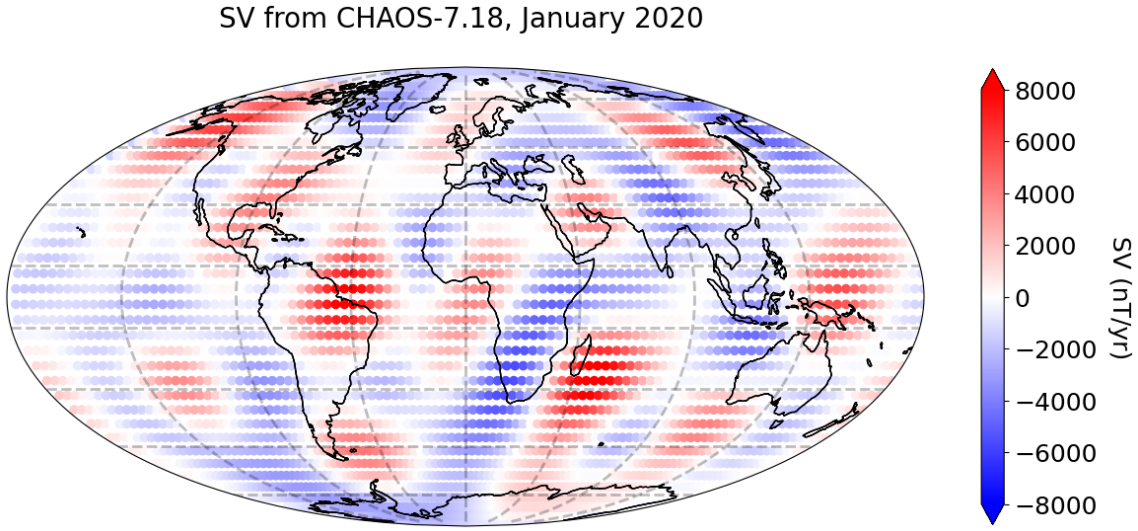



Figure 5.2: SV ($\partial B_r/\partial t$) from CHAOS-7.18 at epoch 2020.0 (Finlay et al. (2020)) at the CMB, truncated at degree 8, with continents shown for reference. Grey lines indicate the edges of the regional latitude-longitude boxes after border removal.

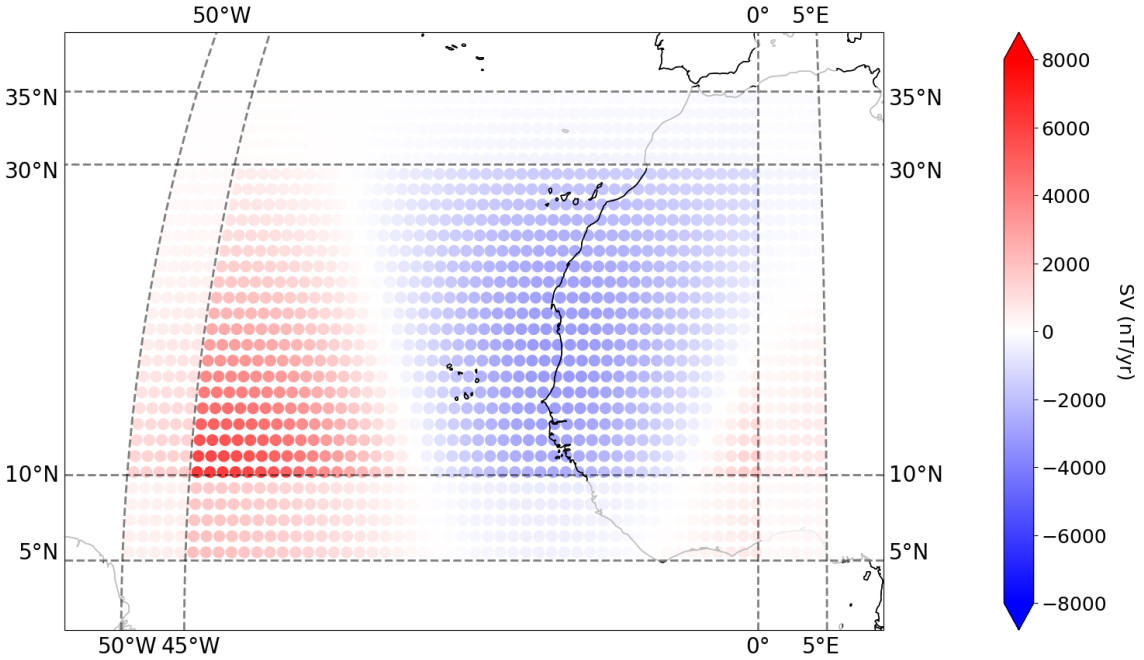



Figure 5.3: Example of a box used for inversion. After training, a 5° border is removed in order to minimise edge effects, this is shown with a transparent overlay. SV ($\partial B_r/\partial t$) from CHAOS-7.18 at $l = 8$, shown for reference. Adapted from Shakespeare-Rees et al. (2025).

## 5.3 Advection Methodology

Once we have found a flow model consistent with the SV for a particular year, we predict the SV assuming advection from a steady flow. We evolve the frozen-flux induction equation over timesteps of 1 year. To do this, we require a flow and an estimate of the main field at the beginning epoch. For each yearly time step $t_i$ through this period, the main field was evolved using:

$$\boldsymbol{B}_{t_i+1} = \boldsymbol{B}_{t_i} + \Delta t \frac{\partial \boldsymbol{B}}{\partial t}\bigg|_{t_i}, \tag{5.7}$$

where the subscripts $t_i$ indicate evaluation at time $i$ and $\Delta t = 1$. Numerically, we evaluate the SV using the PINN machinery on the Gauss-Legendre-Fourier grid, and transform to find the component of spherical harmonic degree 8. The required horizontal field derivatives are found by projecting the field onto spherical harmonics and taking the derivatives in space. No explicit additional flow smoothing, either spatially or temporally, are used in this SV prediction. The mean of these SV snapshots is then determined by summing together each SV snapshot and then dividing by 5, such that:

$$\left(\frac{\partial B_r}{\partial t}\right)_{Average} = \frac{1}{5}\sum_{i=1}^{5} \frac{\partial B}{\partial t}\bigg|_{t_i}. \tag{5.8}$$

## 5.4 Hindcasting and comparison to IGRF-13

### 5.4.1 Regional and Global Models

To evaluate our methodology against the performance of the IGRF-13 secular variation (SV) prediction candidates, we conducted a 4-year hindcast beginning on 1 January 2020 through to 1 January 2024, and compared the results with time-dependent SV from CHAOS-7.18. Only geomagnetic data from 2020 was used for our hindcasts.

The global Mean Absolute Error (MAE) between the SV from the CHAOS-7.18 model and the SV output from the PINN, both at epoch 2020.0, is 52 nT/year, a relative absolute error of 3.2%. The recovered flows for the regional model are shown in Figure 5.4. The flows show similar large scale features as those found using global methodologies, demonstrating westward flow under the Atlantic, eastward flow under the Pacific, and the presence of the anticyclonic planetary gyre (Pais and Jault, 2008; Gillet et al., 2022; Kloss and Finlay, 2019).

For the global approach, we trained a separate PINN for one latitude-longitude box encompassing the non-polar CMB. Again, the PINN for the box is trained for 5000 iterations, and results obtained for 5 seeds. The model with the lowest $L_{SV}$ (SV MSE) is taken as the 'preferred model' for the global box. The MAE between the SV from the CHAOS-7.18 model and the SV output from the PINN is 15 nT/year, a relative L1 error of 0.8%. The recovered flows for the global model are shown in the bottom panel of Figure 5.4. The flows from the global model are visually smoother than that of the regional model, and have reduced small-scale features. There are also sections of southerly flow under Asia, which is not present in the regional model.

Both the global and regional flows are then used in the advection methodology detailed in Section 5.3, to produce global and regional hindcasts. The absolute difference between the global and the regional hindcasts at the CMB are shown in Figure 5.5. The maximum difference in SV hindcast is -400nT/year, largely concentrated underneath South America, with other areas of increased difference including the northern hemisphere, close to the location of the high latitude jet presented in Livermore et al. (2017). The areas of largest difference correspond to the areas of largest flow difference between the two flow models, which is a result of increased small-scale flow.

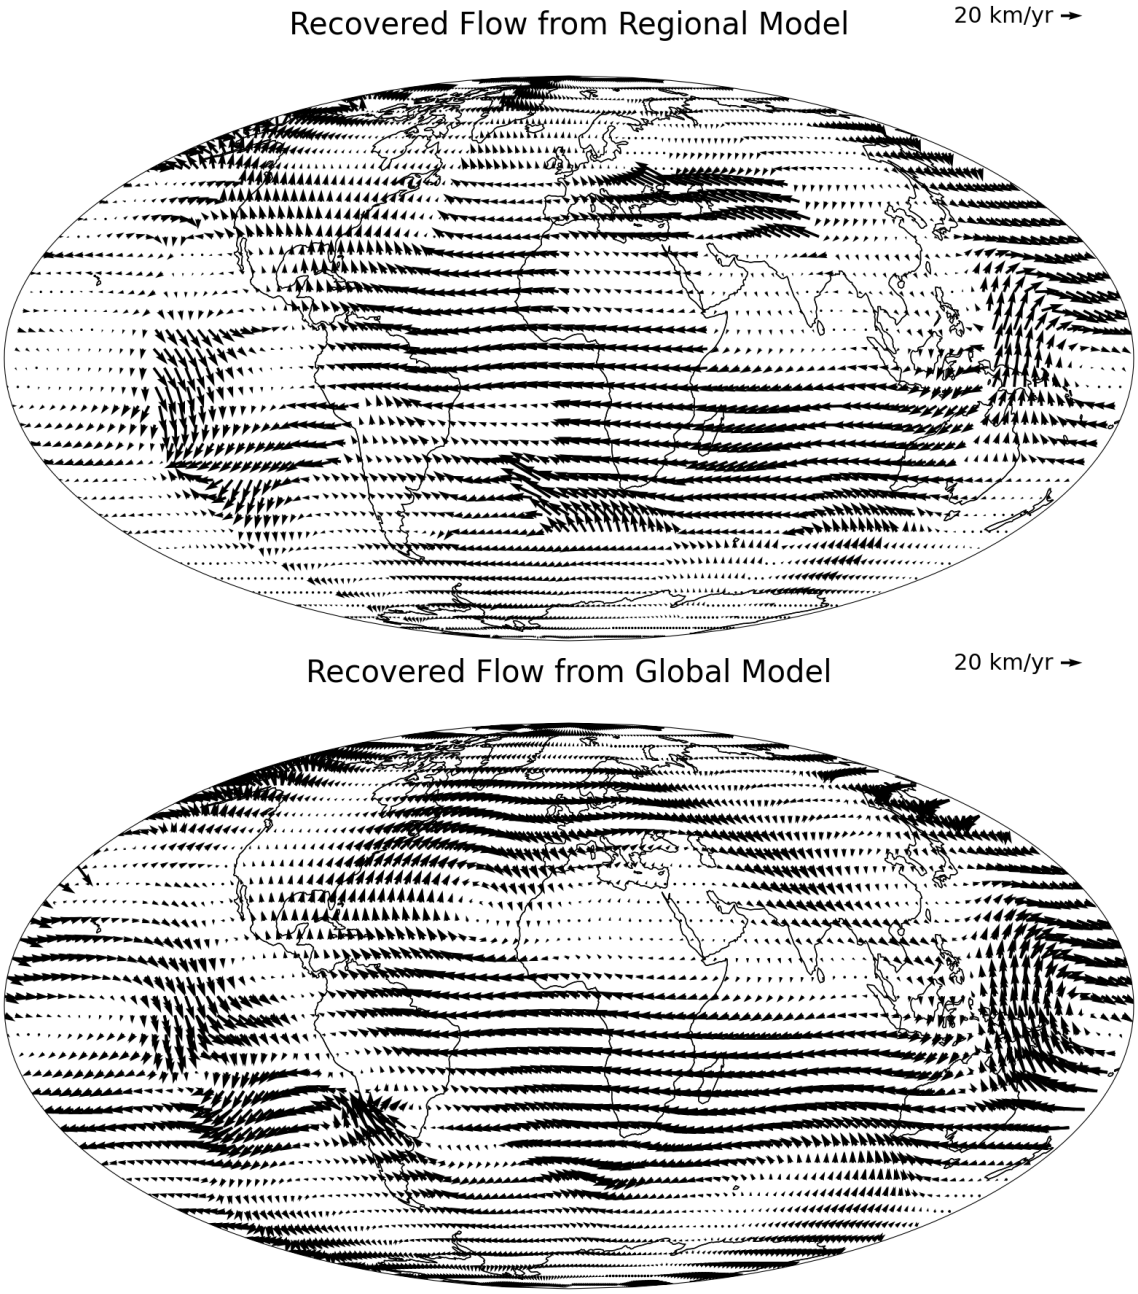



Figure 5.4: Regional (Top) and Global (Bottom) flow models, using the PINNs inversion methodology. All plots are at the CMB. Continents shown for reference.

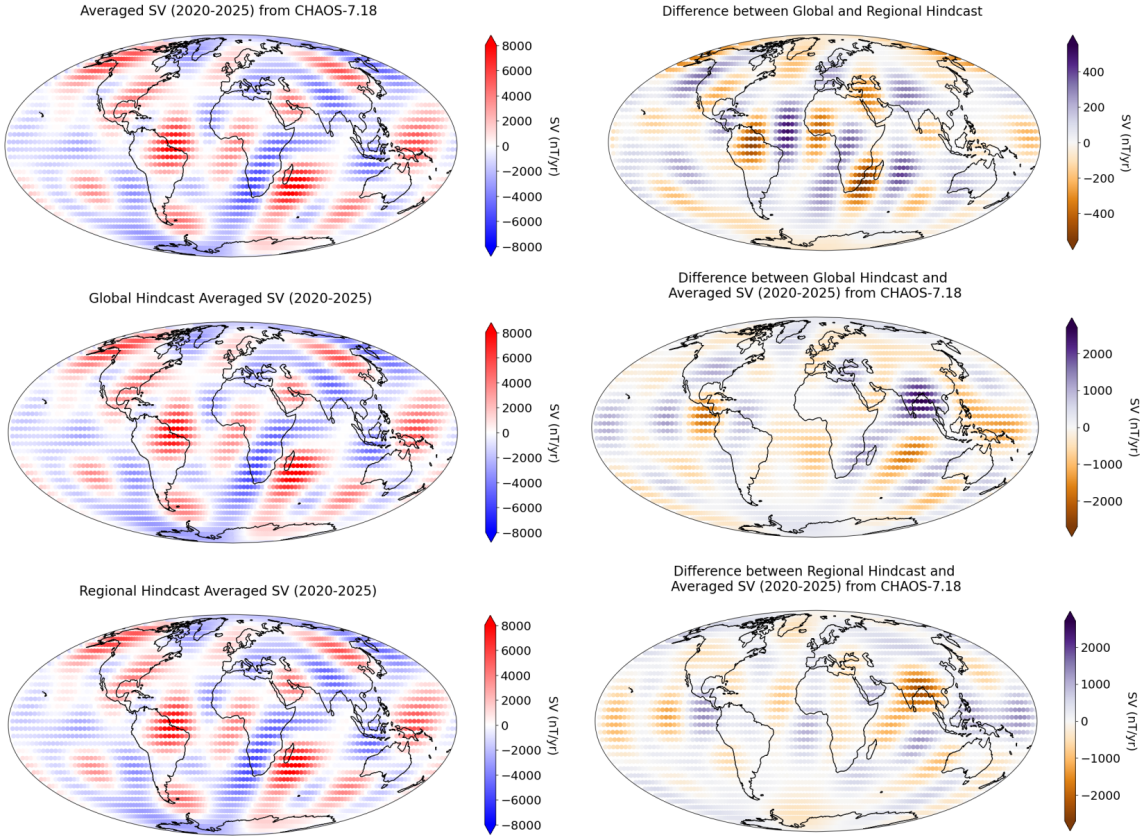



Figure 5.5: Top Left Panel: Average SV ($\partial B_r/\partial t$) over the period 2020–2025 from the CHAOS-7.18 model. Middle Left Panel: Hindcast averaged SV ($\partial B_r/\partial t$) (2020–2025) using the 'Global' flow model. Bottom Left Panel: Hindcast averaged SV ($\partial B_r/\partial t$) (2020–2025) using the 'Regional' flow model. Top Right Panel: Absolute difference between the 'Global' and 'Regional' hindcasts. Middle Right Panel: Absolute difference between the 'Global' hindcast and the average SV ($\partial B_r/\partial t$) over the period 2020–2025 from the CHAOS-7.18 model. Bottom Right Panel: Absolute difference between the 'Regional' hindcast and the average SV ($\partial B_r/\partial t$) over the period 2020–2025 from the CHAOS-7.18 model. All plots are at the CMB. Continents shown for reference.

To further investigate the difference between the global and regional hindcast, the Mausberger-Lowes spectrum of the SV hindcasts was plotted for each model at the Earth's surface. This is shown in Figure 5.6, with the difference between them indicated by the dashed line. These models are almost identical, having a relative L1 difference of 1.8% relative to the regional model at the Earth's surface, indicating that considering one large grid or many smaller grids are very similar, provided the solutions are converged. The relative L1 difference becomes 1.7% relative to the regional model at the CMB.

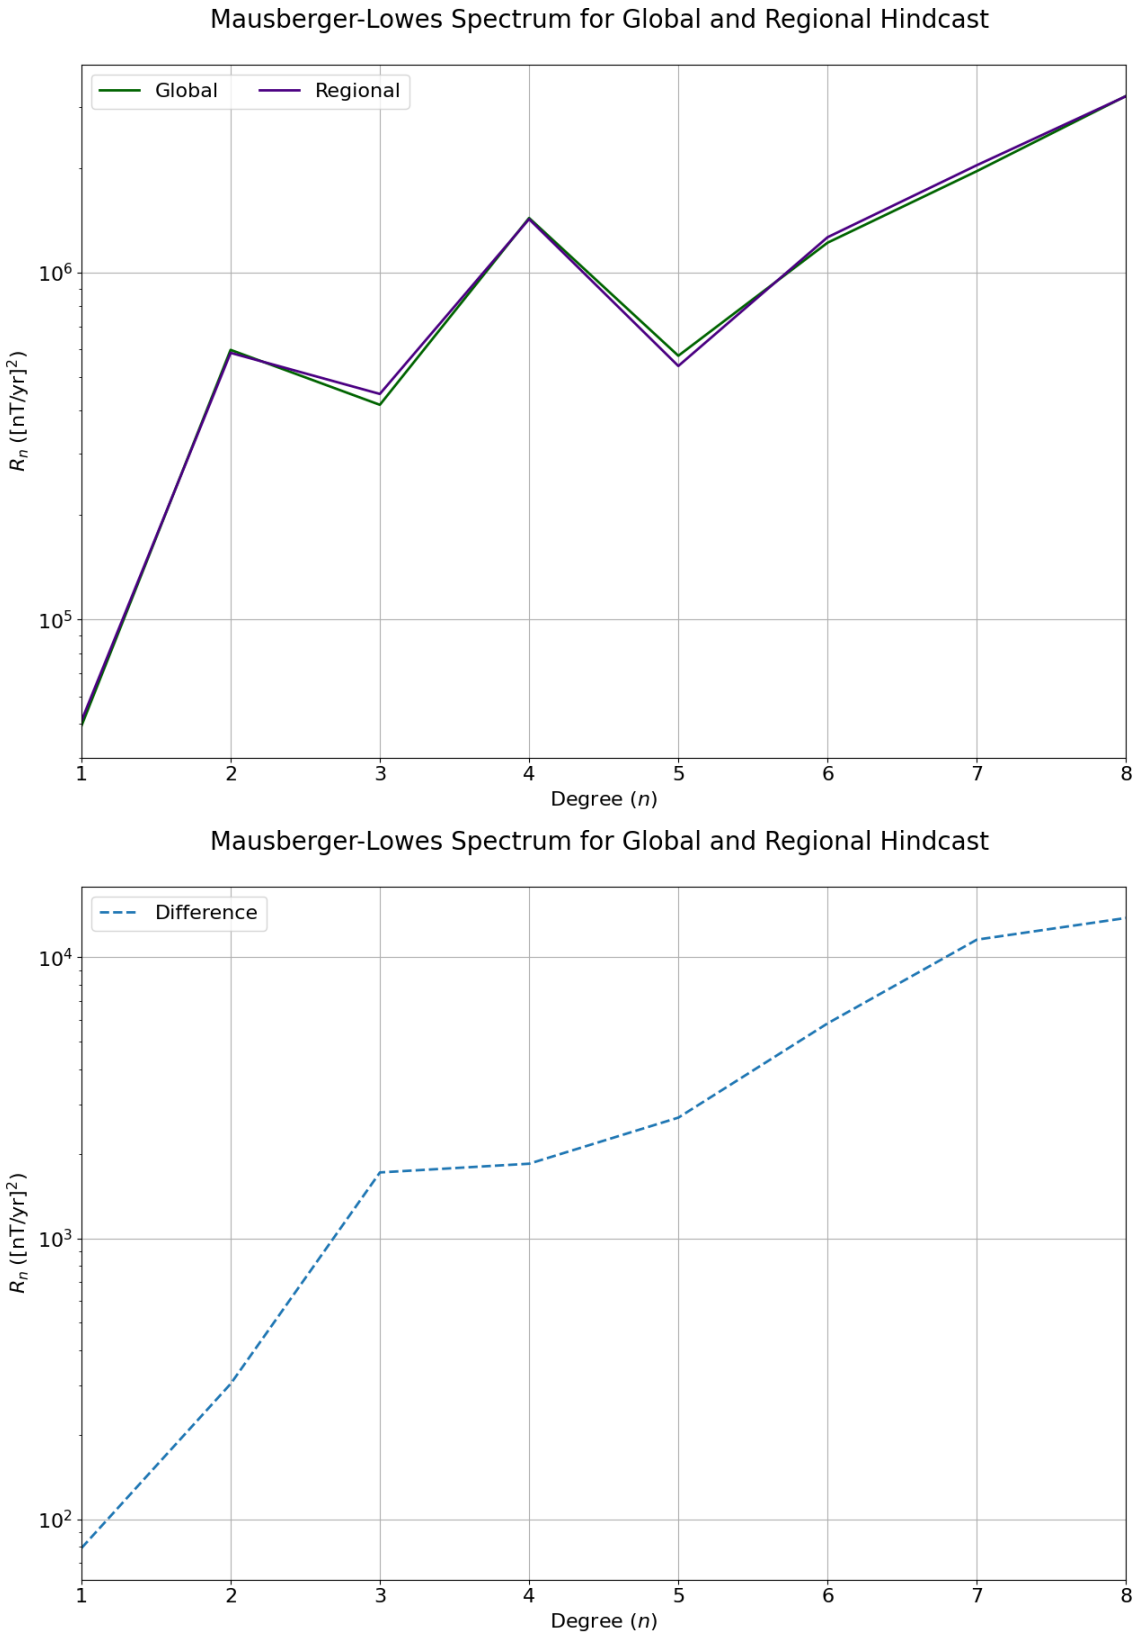



Figure 5.6: Top Panel: Mausberger-Lowes spectrum of the SV for the global and regional hindcasts. Bottom Panel: Difference between the global and regional models. The two spectra are exceedingly similar, having a relative L1 difference of 1.8% relative to the regional model.

### 5.4.2 Comparison to IGRF-13

To compare our hindcasting models to the IGRF-13 SV candidate models, as well as the averaged SV from CHAOS-7.18 between 2020-2024, we first plot the Mausberger-Lowes Spectrum of the SV at the Earth's surface for all of the considered models. The candidate models of IGRF-13 are displayed using dashed lines, the global and regional models presented in this study are displayed using purple and dark green respectively, and the spectra from the average SV between 2020-2024 is shown with a solid black line.

Figure 5.7 shows that the majority of candidate models overestimate the energy at degrees 1-2 and 6-8 compared to the CHAOS-7.18 model, whereas at degree 3-5 the majority of models underestimate the energy. Conversely, both of our models slightly underestimate the degree 1 contribution to the SV, then overestimate the degree 4-8 contributions, with the regional model contributing more at degree 6 than the global model. It is perhaps expected that the regional model contributes more to the SV at higher degrees than the global model, due to the presence of increased small-scale flow, increasing the complexity of the model. Reassuringly, all of these models are similar to the observed SV, particularly at larger scales, though our models follow the black line more closely than many of the other models.

To further probe the hindcast performance, we evolved the main field between 2020.0-2024.0, from an initial state at 2020 derived from CHAOS-7.18, using each of the candidate average SV predictions and our own, such that for each year, and each candidate:

$$\boldsymbol{B}_{t_i+1} = \boldsymbol{B}_{t_i} + \Delta t \frac{\partial \boldsymbol{B}}{\partial t}\bigg|_{Average}, \tag{5.9}$$

where $\Delta t = 1$. The Root Mean Square (RMS) of the difference between the predicted $B_r$ and the modelled $B_r$ from CHAOS-7.18, at degree 8 and at each timestep, was then calculated for each candidate both at the CMB and the Earth's surface.

Figure 5.8 (top panel) shows that both of our models have reduced RMS compared to the other IGRF candidates, at the CMB, with the regional model performing slightly better than the global model. To ensure the robustness of this result and confirm it was not an artifact of our model choice, we compared the predicted SV from flows trained on CHAOS-7.18 with SV from the Kalmag model (Baerenzung et al., 2020). In this comparison, both the global and regional models continued to show reduced RMS difference compared to the other candidate models.

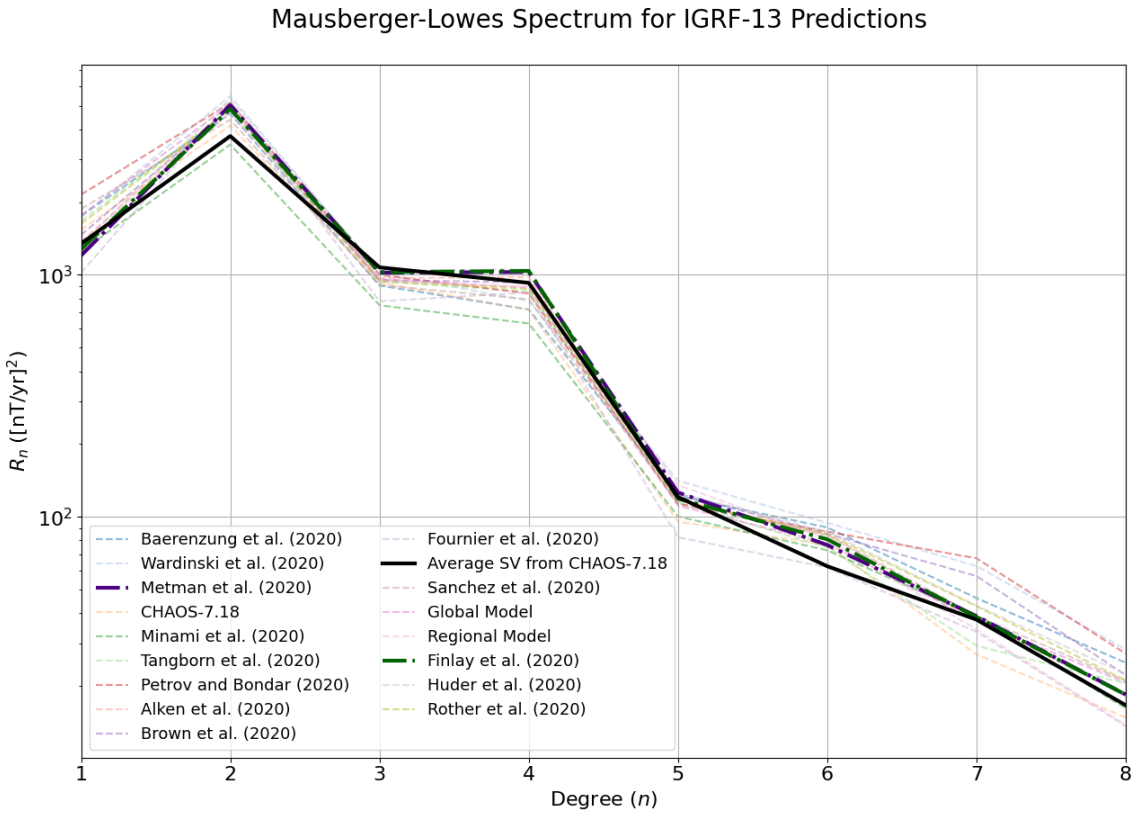



Figure 5.7: Mausberger-Lowes spectrum for our global and regional hindcasts, as well as the other IGRF-13 candidate models. The inferred Mausberger-Lowes Spectrum from the average SV between 2020-2025 from the CHAOS-7.18 model is marked by 'Average SV from CHAOS-7.18'.

This is shown in the figures in Appendix 5.B. The bottom panel of Figure 5.8 shows the results at the Earth's surface, relative to CHAOS-7.18. Once again, our regional model demonstrates reduced RMS differences compared the other models in this hindcast. However, our global model is outperformed by the models from the Max Planck Institute (Sanchez et al., 2020), IPGP (Fournier et al., 2021), Potsdam (Baerenzung et al., 2020) and DTU (Finlay et al., 2020). Therefore, it seems that the small-scale flows present in the regional modelling contribute to the SV in a way that improves forecast performance.

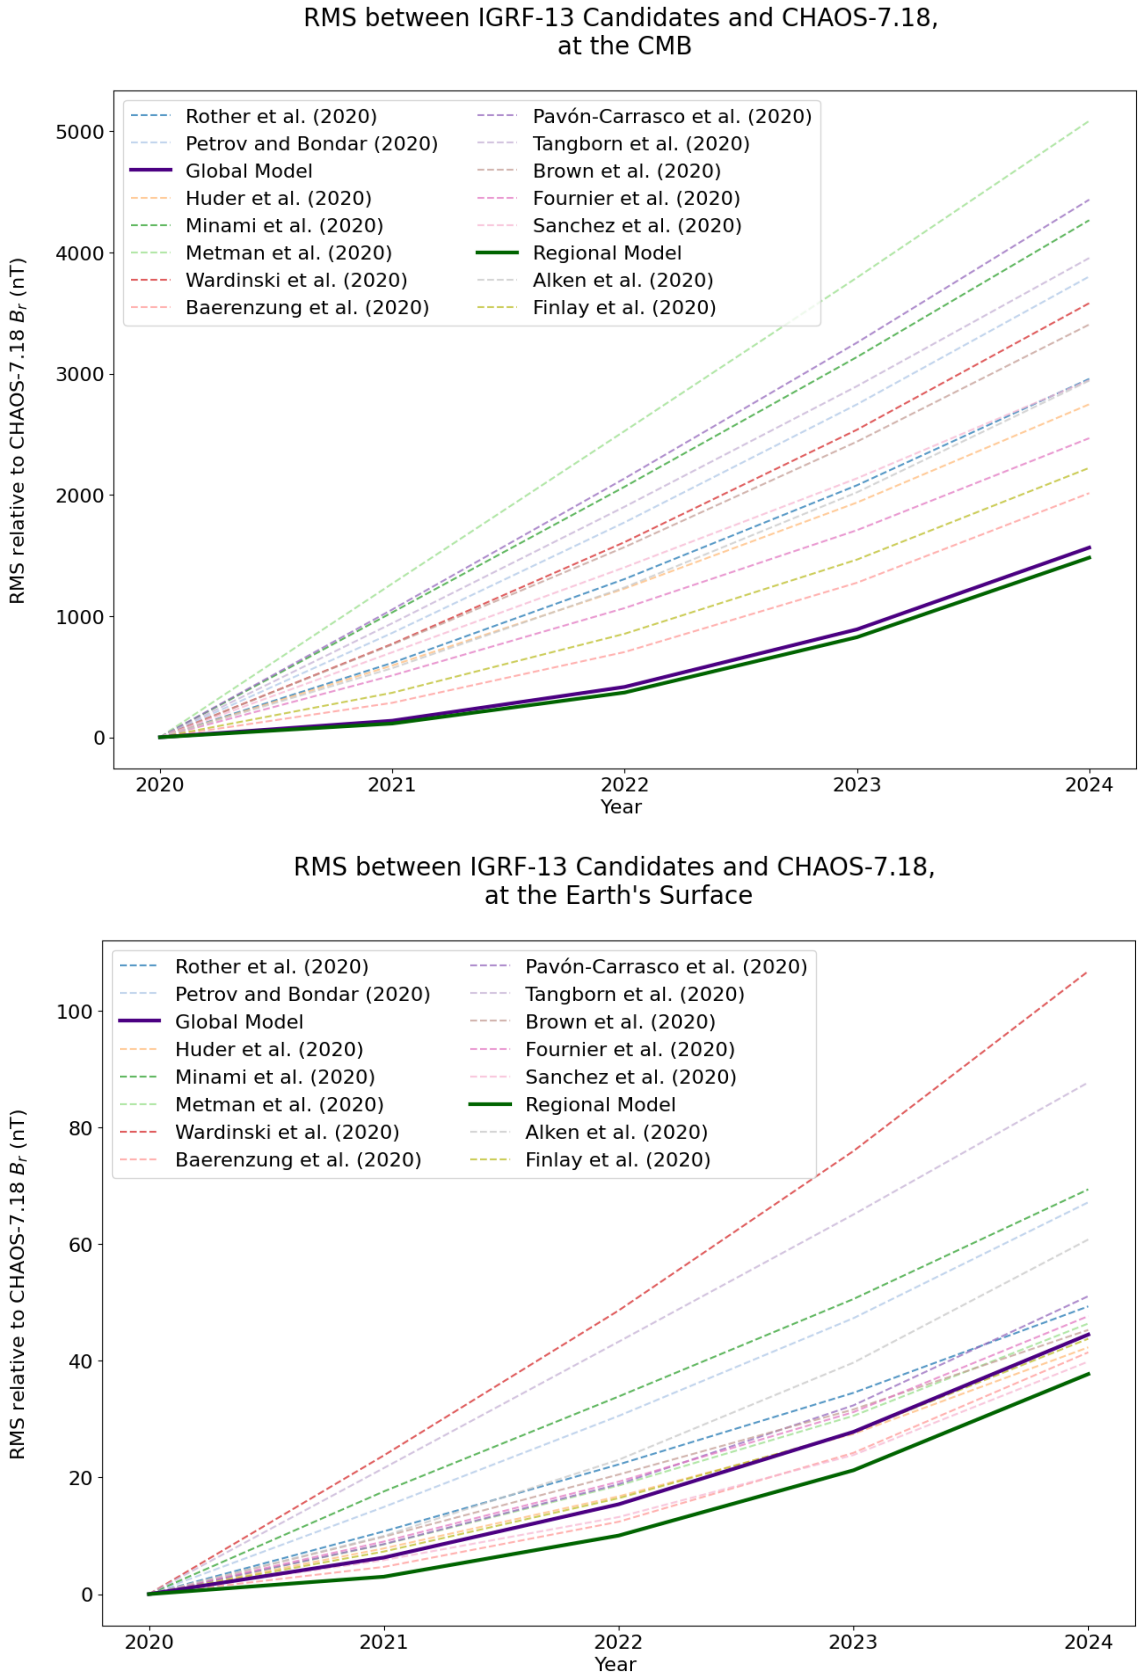



Figure 5.8: RMS between IGRF-13 candidates and CHAOS-7.18 derived $B_r$ values over time. The regional model performs marginally better than the global model, with the difference becoming more pronounced over time.

## 5.5 Candidate SV Model

With the hindcast performance of the regional model in mind, we chose this methodology to determine the flows used in our steady-flow prediction methodology. We used the same methodology as the hindcast, but using the CHAOS-7.18 geomagnetic field data from 1st January 2024. The radial magnetic field at 2025 was estimated by adding the SV on 1st January 2024 to the radial magnetic field at the same date, in the same manner as in Equation (5.7), with $\mathrm{d}t = 1$ year, to predict the field in 2025. Ideally we would use the field in 2025, but this is outside the timeframe of CHAOS-7.18. The recovered SV from the PINN, and the residual between the CHAOS-7.18 SV and the recovered SV, are shown in the top panels of Figure 5.9. At the epoch for which the PINN is trained, the recovered SV agrees well with the training SV, as the MAE between the recovered SV and CHAOS SV is 50 nT/year, a relative L1 error of 3.2%. This is similar to the hindcast values, giving hope that the forecast may fair similarly well. The recovered flows are shown in the bottom panel of Figure 5.9, again showing westward drift under Atlantic, eastward flow under Indonesia, and the eccentric planetary gyre.

We then advect the magnetic field, assuming the flow is steady and using the estimated values of the radial magnetic field at 2025. We note that due to a data transmission error at this step, the original University of Leeds SV prediction candidate for IGRF-14 submitted in October 2024 was incorrect. We show the corrected University of Leeds SV prediction candidate for 2025-2030 in Figure 5.10, and provide the updated Gauss coefficients in Appendix 5.A.

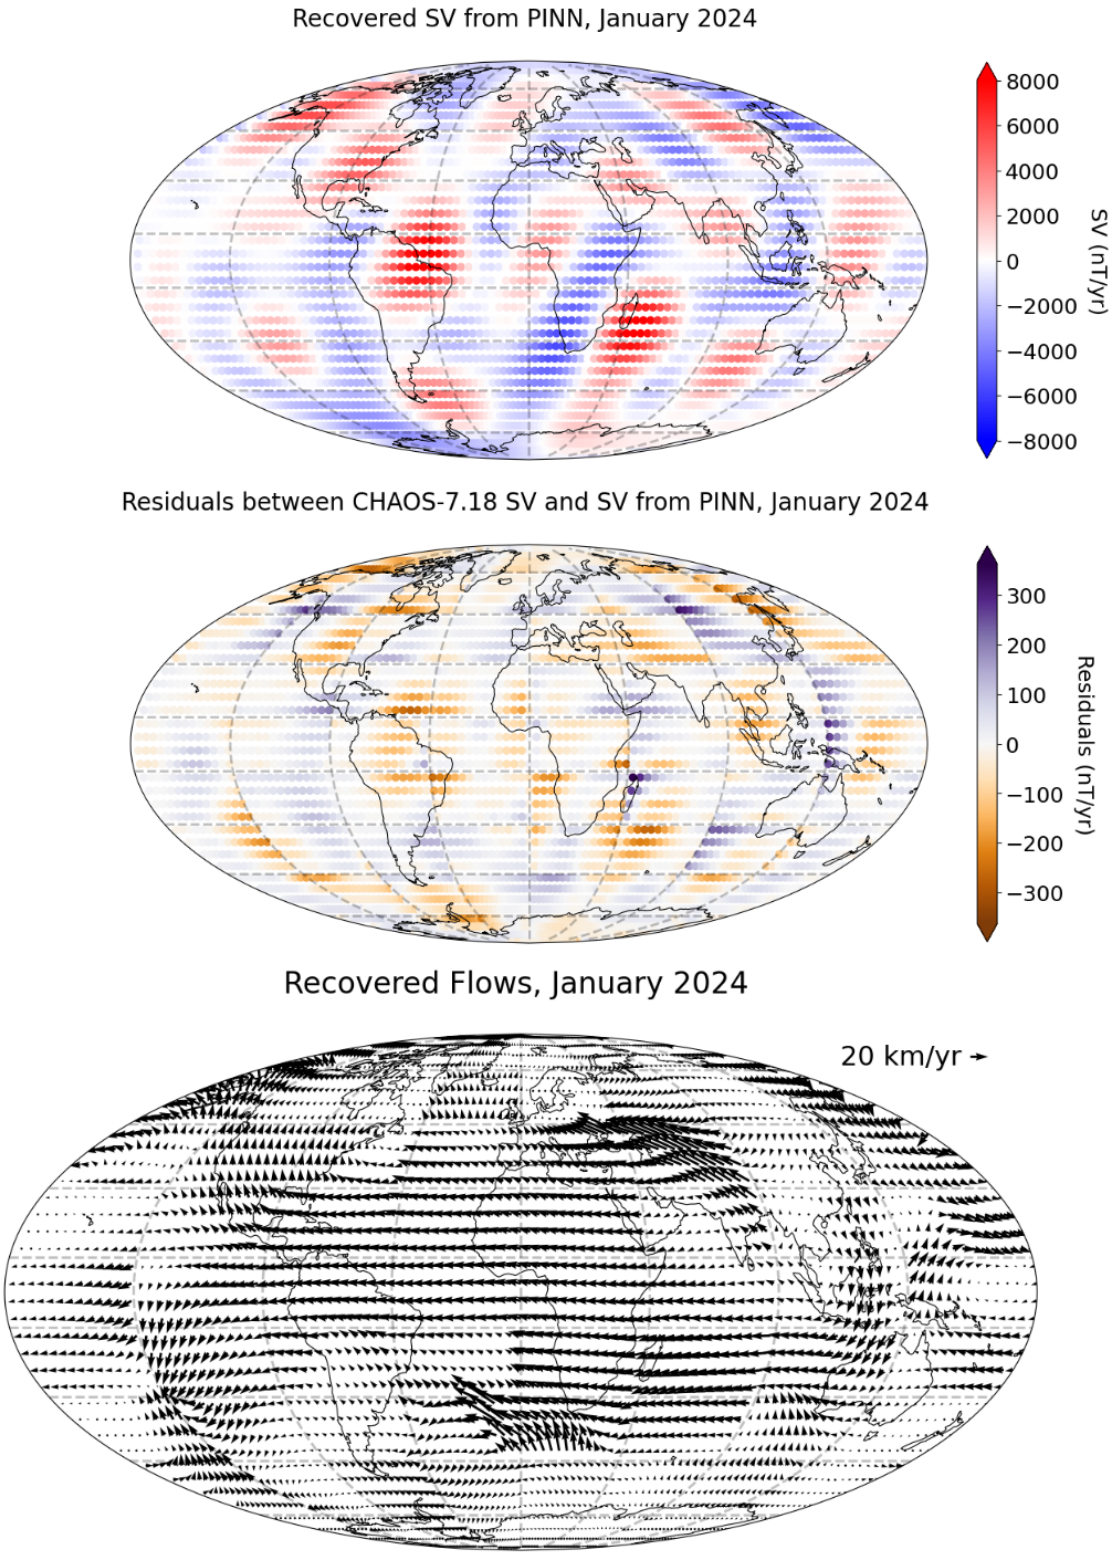



Figure 5.9: Top Panel: Recovered SV using the horizontal flows in January 2024, with the box borders shown using the grey dashed lines. Middle Panel: Residual between CHAOS-7.18 SV and the Recovered SV. Bottom Panel: Recovered Flows from the PINN. The flows show large-scale features found in prior studies, such as westward drift under Atlantic, eastward flow under Indonesia, and the eccentric planetary gyre. Continents shown for reference. The flow at epoch 2024 is assumed to be the same as for 2025 for the purposes of SV prediction.

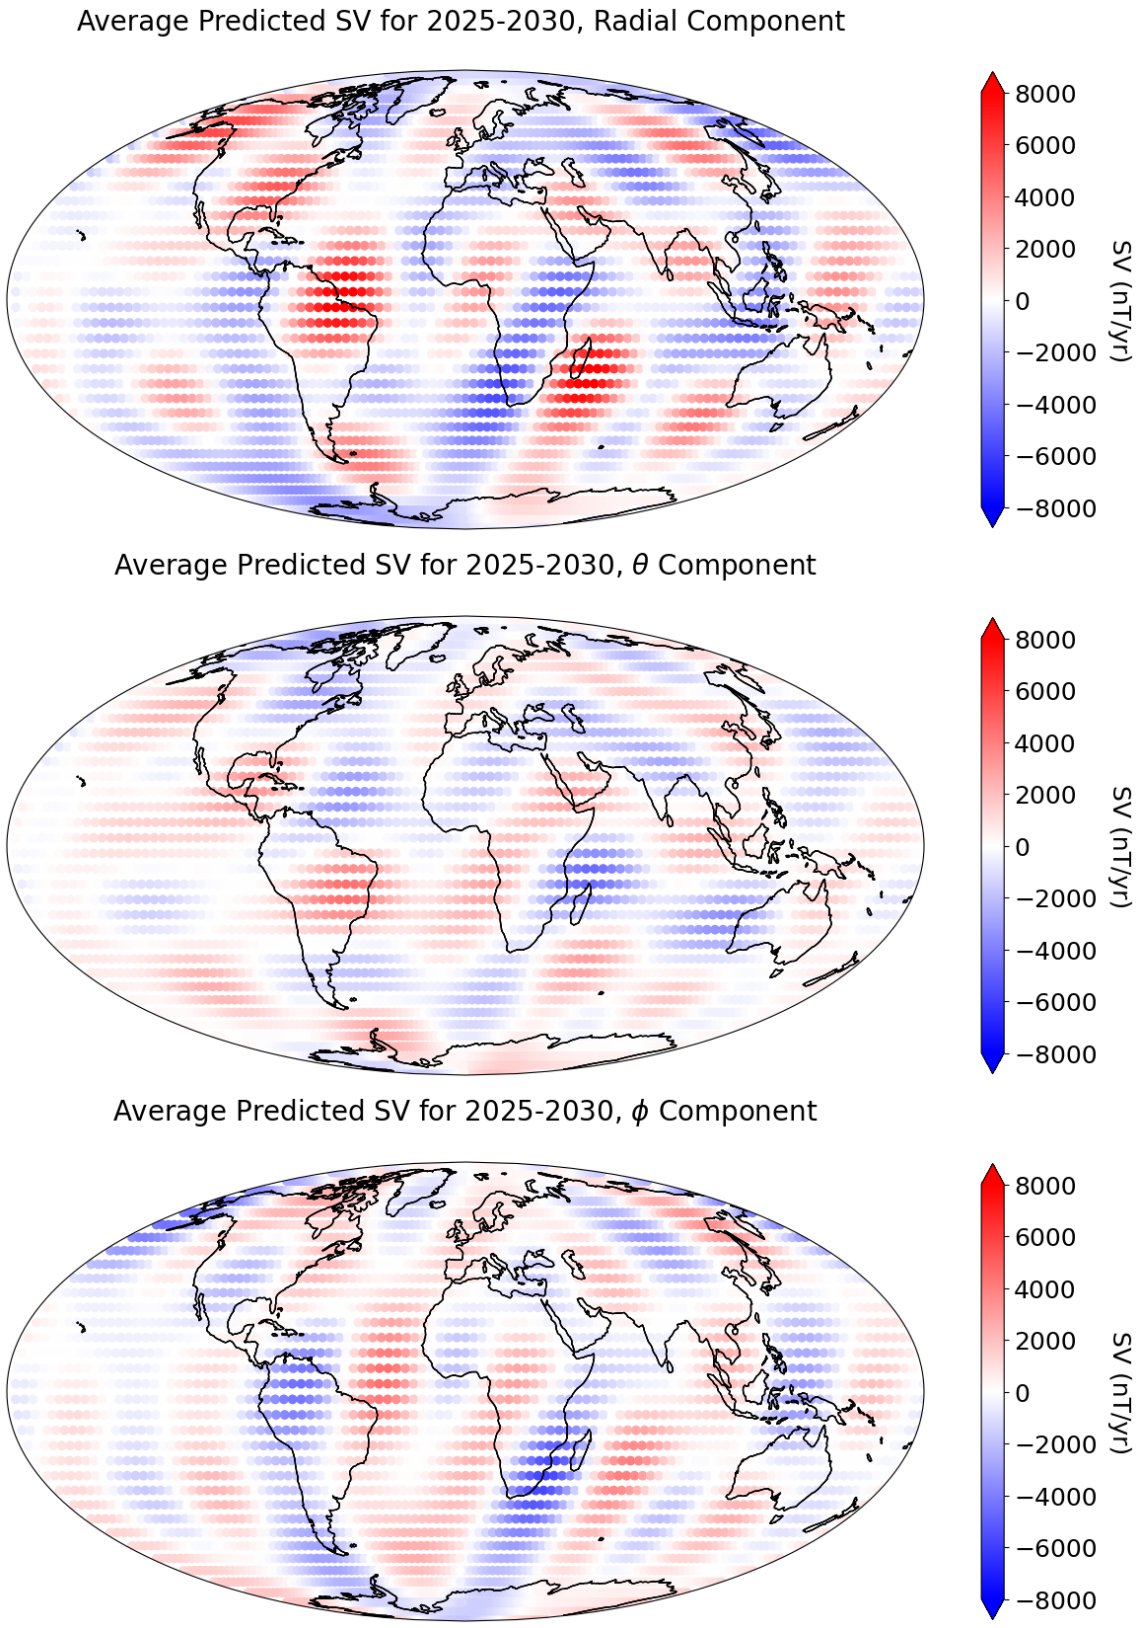



Figure 5.10: Predicted SV (Top: $\partial B_r/\partial t$, Middle: $\partial B_\theta/\partial t$, Bottom: $\partial B_\phi/\partial t$) for epochs 2025.0-2030.0 at the core-mantle boundary. Continents shown for reference.

## 5.6 Discussion and Conclusions

We present a methodology for SV forecasting, using Physics-Informed Neural Networks to determine outer core flows for use in a steady-flow advection method in order to predict SV for IGRF-14. To investigate the performance of the method, we produce a hindcast for the epochs 2020.0-2025.0, using a mosaic of regional flows ('Regional Model') and a global flow ('Global Model').

We then compared our hindcasts to other IGRF-13 SV forecast candidate models, as well as to the average CHAOS-7.18 SV between epochs 2020-2025. It should be noted that the average SV during this time period — what is predicted in our hindcasts and IGRF-13 candidate models — may differ from CHAOS-7.18 SV observations, as un-modelled processes in the outer core leading to variable SV change over time, such as the presence of geomagnetic jerks (Brown et al., 2013; Aubert et al., 2022). The results show that the regional models outperformed the global model and the other IGRF-13 candidates, both at the CMB and at the Earth's surface.

There are a number of considerations and potential improvements that need to be accounted for when evaluating our models. The first is the box size defining both regional and global flows, along with the (removed) additional border. While the border removal approach works well for the regional boxes, acting to mitigate any edge effects, these edge effects are still seen in the global model. Adding a border region did not act to prevent this, and so further work may include altering the global methodology to remove this, as has been done in Li et al. (2024). The use of MSE for in our loss function may also affect this, as when the loss is minimised the optimiser will prioritise a fit over the bulk of the box, rather than the edges which comprise only a small part.

The second consideration is our choice of geomagnetic field model for use as training data. We chose CHAOS-7.18 for the input field model, due to its use in multiple core inversion studies, such as Gillet et al. (2022) and Kloss and Finlay (2019). Other geomagnetic field models such as Kalmag (Baerenzung et al., 2020) also covered the time frame considered in this study, and the use of this model as our training data may have yielded slightly different results. However, the high quality data from INTERMAGNET (Love and Chulliat, 2013) and Swarm mean that these field models are well constrained, such that there is little difference between Kalmag and CHAOS. Since the publication of IGRF-14, the 8th generation of CHAOS has been

released (Kloss et al., 2024). CHAOS-8 is temporally regularised by taking temporal covariances from geodynamo simulations, whereas CHAOS-7 is temporally regularised by penalising the second time derivative of the radial field over the CMB at the endpoints, and its third time derivative throughout. This change in regularisation method influences the small-scale features of the SV, which may, in turn, impact the recovered flow patterns. Further work might include incorporating this model into our prediction methodology, in order to see whether these small scale features increase the forecasting abilities of our model.

Thirdly, while previous studies involving the use of PINNs for core flow inversions have applied the Tangential Geostrophy flow constraint (Li et al., 2024; Shakespeare-Rees et al., 2025), in this modelling approach we chose not to apply a flow constraint, due to the complexities of balancing this constraint in the loss function. Subsequent work may include quantifying the effect of different flow assumptions on forecasting ability.

Finally, we note that the regional models are first trained on grids equivalent to 1 point per spatial degree, and then interpolated to the same Gauss-Legendre-Fourier grid that the global methodology is trained on. Further work will need to be done to investigate whether the enhanced performance of the regional methodology is due to the difference in training grid, or if the regional aspect of the model provides information that aids the forecast.

Our regional Physics-Informed Neural Network approach to forecasting secular variation demonstrates slightly improved performance over previous generations of SV forecasts in hindcasting experiments, though it remains to be seen how our model performs in the next 5 years. As the error between the predicted and 'true' $B_r$ increases with time (Figure 5.8), future work could explore whether weighting earlier time steps more strongly in Equation (5.8) improves forecasting performance. With ongoing advances in machine learning, forecasting methodologies that leverage these techniques will become increasingly prevalent in geomagnetic research. The integration of regional methodologies — such as the one presented in this study — with the additional coverage provided by low-inclination satellite missions like MSS-1 (Jiang et al., 2024) and NanoMagSat (Deconinck et al., 2025) holds particular promise for enhancing the accuracy of future secular variation forecasts, given the similarly regional nature of the data they provide. Further additions to this methodology could include incorporating time-dependent constraints, leveraging the full CHAOS-7.18 time series. Localised estimates, such as those from virtual observatories (Mandea and Olsen, 2006; Hammer et al., 2021) or from SOLA (Subtractive Op-

timally Localised Averages, Hammer and Finlay (2018)), could also be trialled, ensuring there is no global constraint at any stage. Since conventional global field and flow modelling techniques work with purely regional data, the development of robust regional approaches is essential — not only to fully leverage upcoming missions, but also as a necessary redundancy should Swarm fail, or if NanoMagSat launches without its planned polar-orbiting component.

## 5.7 Summary

This study presents a novel methodology to forecast SV, using flows from a Physics-Informed Neural Network based approach in a steady-flow advection scheme. Two approaches were presented: an approach comprising a mosaic of regional flow inversions and a single global flow inversion. Hindcasting results reveal that our regional approach outperforms the IGRF-13 SV candidate models, both at the CMB and the Earth's surface. We present our updated candidate SV model over the time period 2025-2030 for IGRF-14, derived from advected flows from the regional inversion approach.

## 5.A Updated Gauss Coefficients for Leeds' Candidate SV Model

| $n$ | $m$ | $\dot{g}_{nm}$ | $\dot{h}_{nm}$ |
|---|---|---|---|
| 1 | 0 | 14.47 | 0.0 |
| 1 | 1 | 9.38 | -22.10 |
| 2 | 0 | -11.30 | 0.0 |
| 2 | 1 | -5.24 | -27.36 |
| 2 | 2 | -8.75 | -11.15 |
| 3 | 0 | -2.02 | 0.0 |
| 3 | 1 | -4.47 | 3.82 |
| 3 | 2 | 0.69 | -0.19 |
| 3 | 3 | -16.12 | -4.11 |
| 4 | 0 | -1.61 | 0.0 |
| 4 | 1 | -2.48 | -1.21 |
| 4 | 2 | -5.92 | 4.10 |
| 4 | 3 | 5.58 | 1.40 |
| 4 | 4 | -7.02 | -4.0 |
| 5 | 0 | 0.64 | 0.0 |
| 5 | 1 | 1.32 | -0.58 |
| 5 | 2 | 0.14 | 2.17 |
| 5 | 3 | 0.56 | 0.51 |
| 5 | 4 | 2.47 | 1.49 |
| 5 | 5 | 0.79 | 2.13 |
| 6 | 0 | -0.13 | 0.0 |
| 6 | 1 | -0.46 | 0.35 |
| 6 | 2 | 0.85 | -1.50 |
| 6 | 3 | 1.17 | -0.34 |
| 6 | 4 | -0.82 | 0.87 |
| 6 | 5 | 0.41 | 0.72 |
| 6 | 6 | 0.93 | 0.94 |
| 7 | 0 | -0.25 | 0.0 |
| 7 | 1 | -0.06 | 0.64 |
| 7 | 2 | -0.12 | 0.49 |
| 7 | 3 | 0.47 | -0.76 |
| 7 | 4 | -0.12 | 0.11 |
| 7 | 5 | -0.78 | -0.97 |
| 7 | 6 | -0.81 | 0.63 |
| 7 | 7 | 0.92 | -0.25 |
| 8 | 0 | -0.13 | 0.0 |
| 8 | 1 | 0.22 | -0.25 |
| 8 | 2 | 0.04 | 0.45 |
| 8 | 3 | 0.49 | -0.41 |
| 8 | 4 | -0.14 | 0.39 |
| 8 | 5 | 0.25 | -0.57 |
| 8 | 6 | 0.062 | -0.68 |
| 8 | 7 | -0.039 | 0.28 |
| 8 | 8 | 0.22 | 0.23 |

Table 5.1: Updated Gauss coefficients for the University of Leeds IGRF Candidate model. These coefficients can be found in .cof format here: (`https://github.com/geonaomi/LeedsIGRF/blob/main/Leeds_IGRF14_update.cof`).

## 5.B Comparison between IGRF-13 Candidates and Kalmag

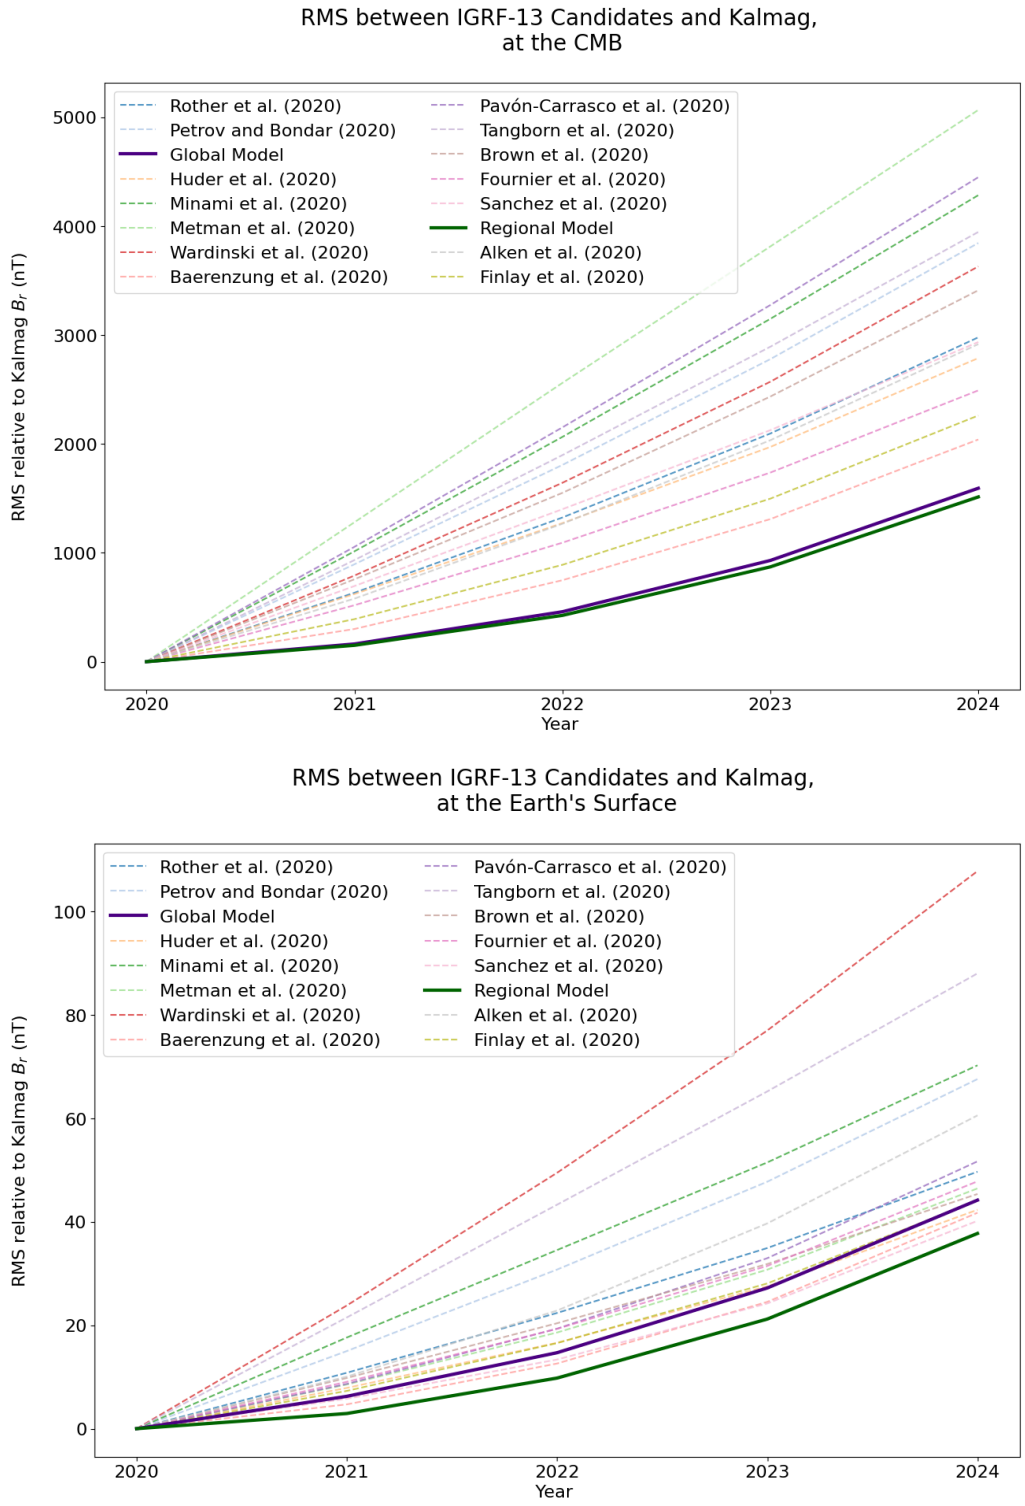



Figure 5.11: RMS between IGRF-13 candidates and Kalmag derived $B_r$ values over time. The regional model performs marginally better than the global model, with the difference becoming more pronounced over time.

## References

Alken, P., Chulliat, A., Nair, M., 2021a. Noaa/ncei and university of colorado candidate models for igrf-13. Earth, Planets and Space 73, 44. URL: https://doi.org/10.1186/s40623-020-01313-z, doi:10.1186/s40623-020-01313-z.

Alken, P., Thébault, E., Beggan, C.D., Amit, H., Aubert, J., Baerenzung, J., Bondar, T.N., Brown, W.J., Califf, S., Chambodut, A., Chulliat, A., Cox, G.A., Finlay, C.C., Fournier, A., Gillet, N., Grayver, A., Hammer, M.D., Holschneider, M., Huder, L., Hulot, G., Jager, T., Kloss, C., Korte, M., Kuang, W., Kuvshinov, A., Langlais, B., Léger, J.M., Lesur, V., Livermore, P.W., Lowes, F.J., Macmillan, S., Magnes, W., Mandea, M., Marsal, S., Matzka, J., Metman, M.C., Minami, T., Morschhauser, A., Mound, J.E., Nair, M., Nakano, S., Olsen, N., Pavón-Carrasco, F.J., Petrov, V.G., Ropp, G., Rother, M., Sabaka, T.J., Sanchez, S., Saturnino, D., Schnepf, N.R., Shen, X., Stolle, C., Tangborn, A., Tøffner-Clausen, L., Toh, H., Torta, J.M., Varner, J., Vervelidou, F., Vigneron, P., Wardinski, I., Wicht, J., Woods, A., Yang, Y., Zeren, Z., Zhou, B., 2021b. International geomagnetic reference field: the thirteenth generation. Earth, Planets and Space 73, 49. doi:10.1186/s40623-020-01288-x.

Arzani, A., Wang, J.X., D'Souza, R.M., 2021. Uncovering near-wall blood flow from sparse data with physics-informed neural networks. Physics of Fluids 33, 071905. doi:10.1063/5.0055600. arXiv:2104.08249 [physics].

Aubert, J., Livermore, P.W., Finlay, C.C., Fournier, A., Gillet, N., 2022. A taxonomy of simulated geomagnetic jerks. Geophysical Journal International 231, 650–672. URL: https://doi.org/10.1093/gji/ggac212, doi:10.1093/gji/ggac212, arXiv:https://academic.oup.com/gji/article-pdf/231/1/650/44394280/ggac212.pdf.

Baerenzung, J., Holschneider, M., Wicht, J., Lesur, V., Sanchez, S., 2020. The kalmag model as a candidate for igrf-13. Earth, Planets and Space 72, 163. URL: https://doi.org/10.1186/s40623-020-01295-y, doi:10.1186/s40623-020-01295-y.

Bloxham, 1988. The dynamical regime of fluid flow at the core surface. Geophysical Research Letters 15, 585–588. doi:10.1029/GL015i006p00585.

Bloxham, Jackson, 1991. Fluid flow near the surface of earth's outer core. Reviews of Geophysics 29, 97–120. doi:10.1029/90RG02470.

Brown, W.J., Beggan, C.D., Cox, G.A., Macmillan, S., 2021. The bgs candidate models for igrf-13 with a retrospective analysis of igrf-12 secular variation forecasts. Earth, Planets and Space 73, 42. URL: https://doi.org/10.1186/s40623-020-01301-3, doi:10.1186/s40623-020-01301-3.

Brown, W.J., Mound, J.E., Livermore, P.W., 2013. Jerks abound: An analysis of geomagnetic observatory data from 1957 to 2008. Physics of the Earth and Planetary Interiors 223, 62–76. URL: https://www.sciencedirect.com/science/article/pii/S0031920113000708, doi:10.1016/j.pepi.2013.06.001.

Bullard, E.C., 1950. The transfer of heat from the core of the earth. Geophysical Journal International 6, 36–41. doi:10.1111/j.1365-246X.1950.tb02971.x.

Chen, Y., de Ridder, S.A.L., Rost, S., Guo, Z., Wu, X., Chen, Y., 2022. Eikonal tomography with physics-informed neural networks: Rayleigh wave phase velocity in the northeastern margin of the tibetan plateau. Geophysical Research Letters 49, e2022GL099053. doi:10.1029/2022GL099053.

Constable, C., Constable, S., 2023. A grand spectrum of the geomagnetic field. Physics of the Earth and Planetary Interiors 344, 107090. URL: https://www.sciencedirect.com/science/article/pii/S0031920123001164, doi:https://doi.org/10.1016/j.pepi.2023.107090.

Deconinck, F., Hulot, G., Léger, J.M., Clausen, L.B.N., Pastena, M., Lejault, J.P., Wilkinson,

J., Caballero, D.M., Tostado, M.V., Coïsson, P., Jager, T., Nieto, P., Jørgensen, J.L., 2025. NanoMagSat status: a 3x16U low-Earth orbit constellation to monitor the Earth magnetic field and the ionospheric environment, in: Petrozzi-Ilstad, M. (Ed.), Small Satellites Systems and Services Symposium (4S 2024), International Society for Optics and Photonics. SPIE. p. 135462K. URL: https://doi.org/10.1117/12.3062471, doi:10.1117/12.3062471.

Finlay, C.C., Kloss, C., Olsen, N., Hammer, M.D., Tøffner-Clausen, L., Grayver, A., Kuvshinov, A., 2020. The chaos-7 geomagnetic field model and observed changes in the south atlantic anomaly. Earth, Planets and Space 72, 156. URL: https://doi.org/10.1186/s40623-020-01252-9, doi:10.1186/s40623-020-01252-9.

Fournier, A., Aubert, J., Lesur, V., Ropp, G., 2021. A secular variation candidate model for igrf-13 based on swarm data and ensemble inverse geodynamo modelling. Earth, Planets and Space 73, 43. URL: https://doi.org/10.1186/s40623-020-01309-9, doi:10.1186/s40623-020-01309-9.

Friis-Christensen, E., Lühr, H., Hulot, G., 2006. Swarm: A constellation to study the earth's magnetic field. Earth, Planets and Space 58, 351–358. doi:10.1186/BF03351933.

Gillet, Gerick, F., Jault, D., Schwaiger, T., Aubert, J., Istas, M., 2022. Satellite magnetic data reveal interannual waves in earth's core. Proceedings of the National Academy of Sciences of the United States of America 119, e2115258119. doi:10.1073/pnas.2115258119.

Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. MIT Press. http://www.deeplearningbook.org.

Haghighat, E., Juanes, R., 2021. Sciann: A keras/tensorflow wrapper for scientific computations and physics-informed deep learning using artificial neural networks. Computer Methods in Applied Mechanics and Engineering 373, 113552. doi:10.1016/j.cma.2020.113552. arXiv:2005.08803 [cs].

Hamilton, B., Ridley, V.A., Beggan, C.D., Macmillan, S., 2015. The bgs magnetic field candidate models for the 12th generation igrf. Earth, Planets and Space 67, 69. URL: https://doi.org/10.1186/s40623-015-0227-x, doi:10.1186/s40623-015-0227-x.

Hammer, M.D., Cox, G.A., Brown, W.J., Beggan, C.D., Finlay, C.C., 2021. Geomagnetic virtual observatories: monitoring geomagnetic secular variation with the swarm satellites. Earth, Planets and Space 73, 54. doi:10.1186/s40623-021-01357-9.

Hammer, M.D., Finlay, C.C., 2018. Local averages of the core–mantle boundary magnetic field from satellite observations. Geophysical Journal International 216, 1901–1918. URL: https://doi.org/10.1093/gji/ggy515, doi:10.1093/gji/ggy515, arXiv:https://academic.oup.com/gji/article-pdf/216/3/1901/27493010/ggy515.pdf.

Holme, R., 2007. 8.04 - Large-Scale Flow in the Core. Elsevier, Amsterdam. p. 107–130. URL: https://www.sciencedirect.com/science/article/pii/B9780444527486001279, doi:10.1016/B978-044452748-6.00127-9.

Holme, R., 2015. 8.04 - large-scale flow in the core, in: Schubert, G. (Ed.), Treatise on Geophysics (Second Edition). second edition ed.. Elsevier, Oxford, pp. 91–113. URL: https://www.sciencedirect.com/science/article/pii/B978044453802400138X, doi:https://doi.org/10.1016/B978-0-444-53802-4.00138-X.

Huder, L., Gillet, N., Finlay, C.C., Hammer, M.D., Tchoungui, H., 2020. Cov-obs.x2: 180 years of geomagnetic field evolution from ground-based and satellite observations. Earth, Planets and Space 72, 160. URL: https://doi.org/10.1186/s40623-020-01194-2, doi:10.1186/s40623-020-01194-2.

Jackson, A., Finlay, C., 2015. 5.05 - geomagnetic secular variation and its applications to the core, in: Schubert, G. (Ed.), Treatise on Geophysics (Second Edition). second edition ed.. Elsevier, Oxford, pp. 137–184. URL: https:

//www.sciencedirect.com/science/article/pii/B9780444538024000993, doi:https://doi.org/10.1016/B978-0-444-53802-4.00099-3.

Jiang, Y., Finlay, C.C., Olsen, N., Tøffner-Clausen, L., Yan, Q., Zhang, K., 2024. Macau scientific satellite-1 initial magnetic field model. Geophysical Research Letters 51, e2024GL112305. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2024GL112305, doi:https://doi.org/10.1029/2024GL112305, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2024GL112305. e2024GL112305 2024GL112305.

Jordan, Mitchell, 2015. Machine learning: Trends, perspectives, and prospects. Science 349, 255–260. doi:10.1126/science.aaa8415.

Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization URL: http://arxiv.org/abs/1412.6980, doi:10.48550/arXiv.1412.6980. arXiv:1412.6980 [cs].

Kloss, Finlay, 2019. Time-dependent low-latitude core flow and geomagnetic field acceleration pulses. Geophysical Journal International 217, 140–168. doi:10.1093/gji/ggy545.

Kloss, C., Finlay, C., Olsen, N., Tøffner-Clausen, L., Gillet, N., Grayver, A., 2024. The dtu geomagnetic field parent model and derived igrf-14 candidate models. URL: http://www.spacecenter.dk/files/magnetic-models/CHAOS-8/_downloads/1327fd0859c41d20bd663446966758ed/IGRF14_DTU_Technical_Note.pdf. technical Note for CHAOS-8.1, accessed 7 Jan 2025.

Li, J., Lin, Y., Zhang, K., 2024. Dynamic mode decomposition of the core surface flow inverted from geomagnetic field models. Geophysical Research Letters 51, e2023GL106362. URL: https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2023GL106362, doi:10.1029/2023GL106362.

Livermore, P.W., Hollerbach, R., Finlay, C.C., 2017. An accelerating high-latitude jet in earth's core. Nature Geoscience 10, 62–68. doi:10.1038/ngeo2859.

Love, J.J., Chulliat, A., 2013. An international network of magnetic observatories. Eos, Transactions American Geophysical Union 94, 373–374. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2013EO420001, doi:https://doi.org/10.1002/2013EO420001, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2013EO420001.

Macmillan, S., 2007. Observatories, Overview. Springer Netherlands, Dordrecht. pp. 708–711. URL: https://doi.org/10.1007/978-1-4020-4423-6_224, doi:10.1007/978-1-4020-4423-6_224.

Mandea, Olsen, 2006. A new approach to directly determine the secular variation from magnetic satellite observations. Geophysical Research Letters 33. URL: https://onlinelibrary.wiley.com/doi/abs/10.1029/2006GL026616, doi:10.1029/2006GL026616.

Metman, M.C., Beggan, C.D., Livermore, P.W., Mound, J.E., 2020a. Forecasting yearly geomagnetic variation through sequential estimation of core flow and magnetic diffusion. Earth, Planets and Space 72, 149. URL: https://doi.org/10.1186/s40623-020-01193-3, doi:10.1186/s40623-020-01193-3.

Metman, M.C., Beggan, C.D., Livermore, P.W., Mound, J.E., 2020b. Forecasting yearly geomagnetic variation through sequential estimation of core flow and magnetic diffusion. Earth, Planets and Space 72, 149. URL: https://doi.org/10.1186/s40623-020-01193-3, doi:10.1186/s40623-020-01193-3.

Minami, T., Nakano, S., Lesur, V., Takahashi, F., Matsushima, M., Shimizu, H., Nakashima, R., Taniguchi, H., Toh, H., 2020. A candidate secular variation model for igrf-13 based on mhd dynamo simulation and 4denvar data assimilation. Earth, Planets and Space 72, 136. URL: https://doi.org/10.1186/s40623-020-01253-8, doi:10.1186/s40623-020-01253-8.

Olsen, N., Albini, G., Bouffard, J., Parrinello, T., Tøffner-Clausen, L., 2020. Magnetic observations from cryosat-2: calibration and processing of satellite platform magnetometer data. Earth, Planets and Space 72, 48. URL: https://doi.org/10.1186/s40623-020-01171-9, doi:10.1186/s40623-020-01171-9.

Olsen, N., Floberghagen, R., 2018. Exploring geospace from space: the swarm satellite constellation mission. Space Research Today 203, 61–71. URL: https://www.sciencedirect.com/science/article/pii/S1752929818300884, doi:https://doi.org/10.1016/j.srt.2018.11.017.

Olsen, N., Kotsiaros, S., 2011. Magnetic Satellite Missions and Data. Springer Netherlands, Dordrecht. pp. 27–44. URL: https://doi.org/10.1007/978-90-481-9858-0_2, doi:10.1007/978-90-481-9858-0_2.

Pais, M.A., Jault, D., 2008. Quasi-geostrophic flows responsible for the secular variation of the earth's magnetic field. Geophysical Journal International 173, 421–443. URL: https://doi.org/10.1111/j.1365-246X.2008.03741.x, doi:10.1111/j.1365-246X.2008.03741.x, arXiv:https://academic.oup.com/gji/article-pdf/173/2/421/6078007/173-2-421.pdf.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A., 2017. Automatic differentiation in pytorch. URL: https://api.semanticscholar.org/CorpusID:40027675.

Pavón-Carrasco, F.J., Marsal, S., Torta, J.M., Catalán, M., Martín-Hernández, F., Tordesillas, J.M., 2020. Bootstrapping swarm and observatory data to generate candidates for the dgrf and igrf-13. Earth, Planets and Space 72, 152. URL: https://doi.org/10.1186/s40623-020-01198-y, doi:10.1186/s40623-020-01198-y.

Petrov, V.G., Bondar, T.N., 2021. Izmiran candidate field model for igrf-13. Earth, Planets and Space 73, 46. URL: https://doi.org/10.1186/s40623-020-01312-0, doi:10.1186/s40623-020-01312-0.

Raissi, Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378, 686–707. doi:10.1016/j.jcp.2018.10.045.

Raissi, Yazdani, A., Karniadakis, G.E., 2020. Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. Science (New York, N.Y.) 367, 1026–1030. doi:10.1126/science.aaw4741.

Ramachandran, P., Zoph, B., Le, Q.V., 2017. Swish: a self-gated activation function URL: http://arxiv.org/abs/1710.05941. arXiv:1710.05941 [cs].

Roberts, P.H., Scott, S., 1965. On analysis of the secular variation. Journal of geomagnetism and geoelectricity 17, 137–151. doi:10.5636/jgg.17.137.

Rogers, H.F., 2022. Improving inversions of regional outer core surface flow models through the application of spherical slepian function URL: https://era.ed.ac.uk/handle/1842/39295, doi:10.7488/era/2546. accepted: 2022-08-02T10:34:07Z.

Rother, M., Korte, M., Morschhauser, A., Vervelidou, F., Matzka, J., Stolle, C., 2021. The mag.num core field model as a parent for igrf-13, and the recent evolution of the south atlantic anomaly. Earth, Planets and Space 73, 50. URL: https://doi.org/10.1186/s40623-020-01277-0, doi:10.1186/s40623-020-01277-0.

Sanchez, S., Wicht, J., Bärenzung, J., 2020. Predictions of the geomagnetic secular variation based on the ensemble sequential assimilation of geomagnetic field models by dynamo simulations. Earth, Planets and Space 72, 157. URL: https://doi.org/10.1186/s40623-020-01279-y, doi:10.1186/s40623-020-01279-y.

Schiassi, E., De Florio, M., Ganapol, B.D., Picca, P., Furfaro, R., 2022. Physics-informed neural

networks for the point kinetics equations for nuclear reactor dynamics. Annals of Nuclear Energy 167, 108833. doi:10.1016/j.anucene.2021.108833.

Schwaiger, T., Jault, D., Gillet, N., Schaeffer, N., Mandea, M., 2023. Local estimation of quasi-geostrophic flows in earth's core. Geophysical Journal International 234, 494–511. URL: https://doi.org/10.1093/gji/ggad089, doi:10.1093/gji/ggad089, arXiv:https://academic.oup.com/gji/article-pdf/234/1/494/50509990/ggad089.pdf.

Shakespeare-Rees, N., Livermore, P.W., Davies, C.J., Rogers, H.F., Beggan, C.D., Brown, W.J., Finlay, C.C., 2026. Forecasting secular variation using physics-informed neural networks for IGRF-14. Earth, Planets and Space 78, 98. URL: https://doi.org/10.1186/s40623-026-02427-6, doi:10.1186/s40623-026-02427-6.

Shakespeare-Rees, N., Livermore, P.W., Davies, C.J., Rogers, H.F., Brown, W.J., Beggan, C.D., Finlay, C.C., 2025. Local flow estimation at the top of the earth's core using physics informed neural networks. Physics of the Earth and Planetary Interiors 367, 107424. URL: https://www.sciencedirect.com/science/article/pii/S0031920125001189, doi:https://doi.org/10.1016/j.pepi.2025.107424.

Sola, J., Sevilla, J., 1997. Importance of input data normalization for the application of neural networks to complex industrial problems. IEEE Transactions on Nuclear Science 44, 1464–1468. doi:10.1109/23.589532.

Tangborn, A., Kuang, W., Sabaka, T.J., Yi, C., 2021. Geomagnetic secular variation forecast using the nasa gems ensemble kalman filter: A candidate sv model for igrf-13. Earth, Planets and Space 73, 47. URL: https://doi.org/10.1186/s40623-020-01324-w, doi:10.1186/s40623-020-01324-w.

Wardinski, I., Saturnino, D., Amit, H., Chambodut, A., Langlais, B., Mandea, M., Thébault, E., 2020. Geomagnetic core field models and secular variation forecasts for the 13th international geomagnetic reference field (igrf-13). Earth, Planets and Space 72, 155. URL: https://doi.org/10.1186/s40623-020-01254-7, doi:10.1186/s40623-020-01254-7.

Whaler, K.A., Hammer, M.D., Finlay, C.C., Olsen, N., 2022. Core surface flow changes associated with the 2017 pacific geomagnetic jerk. Geophysical Research Letters 49, e2022GL098616. doi:10.1029/2022GL098616.

Zhang, K., 2023. A novel geomagnetic satellite constellation: Science and applications. Earth and Planetary Physics 7, 4–21. URL: https://www.eppcgs.org/en/article/doi/10.26464/epp2023019, doi:10.26464/epp2023019.

# Chapter 6

# Using autoencoders to classify geodynamo heat flux heterogeneity from snapshots of the magnetic field

Naomi Shakespeare-Rees, Philip W. Livermore, Christopher J. Davies, Hannah F. Rogers

## Abstract

Simulations of the geodynamo demonstrate that heterogeneous heat flow across the core–mantle boundary can influence the morphology and evolution of the geomagnetic field. However, the amplitude and structure of this heterogeneity within the Earth remain poorly constrained. We present a convolutional autoencoder classifier to test whether snapshots of the radial magnetic field at the CMB can be used to infer the imposed class of heat-flux heterogeneity, $Q^*$. By combining reconstruction and classification objectives, the model learns a compressed representation of the magnetic field that can be used to predict the heat-flux class of unseen snapshots. We first apply the method to the simulations of Mound and Davies (2023), which comprise three classes: homogeneous ($Q^* = 0$), weakly heterogeneous ($Q^* = 2.46$) and strongly heterogeneous ($Q^* = 5.35$). The classifier distinguishes these classes with an error of 4.9%. Sensitivity tests show that classification depends on non-dipolar structure and relative field amplitude, with performance decreasing when the field is truncated below spherical harmonic degree $n = 4$. Regional tests and saliency maps indicate that the most useful information is concentrated at equatorial and mid-latitudes, particularly near the Pacific and African Large Low Velocity Provinces (LLVPs). We then apply the method to a wider simulation set with a magnetic-to-kinetic energy ratio $M_{\text{ratio}} \geq 1$, $400 \leq R_m \leq 2200$ and $0.13 \leq Q^* \leq 10.72$, in which the heat-flux patterns and dynamo parameters also vary. The classification error increases to 20%, but decreases to 13% when the weakest heterogeneous cases are removed, suggesting that low-amplitude heterogeneity is difficult to distinguish from homogeneous cases and that the influence of $Q^*$ is harder to isolate when other simulation properties vary. Finally, we apply the classifier to models of the Earth's magnetic field. CHAOS-8.5 is consistently classified as corresponding to homogeneous CMB heat flux, whereas the gufm1 classifications vary between epochs and model initialisations. Overall, the results show that simulated CMB magnetic fields contain classifiable information about imposed heat-flux heterogeneity. However, the contrasting results for CHAOS-8.5 and gufm1 indicate that a broader simulation training set is required before classifications of the Earth's magnetic field can be considered reliable.

## 6.1 Introduction

The Earth's magnetic field is generated by a self-sustaining dynamo driven by fluid motion in the liquid outer core, with the required energy supplied by heat loss from the core to the mantle (Landeau et al., 2022; Frost et al., 2022). Because of its influence on the core flow, and therefore indirectly on the magnetic field, thermal coupling between the core and mantle has been proposed as a mechanism capable of influencing a wide range of geomagnetic behaviour. Heterogeneous heat flux imposed at the Core-Mantle Boundary (CMB) has been used to explain persistent or quasi-stationary magnetic flux patches and aspects of secular variation (Bloxham and Gubbins, 1987; Amit et al., 2010), changes in reversal frequency (Glatzmaier et al., 1999; Frasson et al., 2025), and large-scale non-zonal structure in the geomagnetic field (Mound and Davies, 2023). Related modelling studies have considered whether core–mantle thermal coupling can influence the location of surface magnetic intensity minima associated with the South Atlantic Anomaly (Terra-Nova and Amit, 2025). The redistribution of CMB heat-flux patterns by true polar wander has also been examined, as the resulting changes in the latitudinal distribution of heat flux may affect the stability of the magnetic dipole (Frasson et al., 2024).

Seismic studies reveal heterogeneous structure in the lowermost mantle. The most prominent structures associated with long-wavelength CMB heat-flux anomalies are the large low-velocity provinces (LLVPs). LLVPs are two broad, approximately antipodal structures beneath Africa and the Pacific, covering around 25% of the core surface, and are characterised by reduced seismic velocities (Koelemeijer, 2021). Heat-flux models can be constructed from tomography models such as that of Masters et al. (1996) by assuming that seismic velocity anomalies are related to temperature anomalies. However, evidence that LLVPs are chemically distinct (Zhao et al., 2015), together with the difficulty of separating thermal and chemical effects in tomographic models (Mosca et al., 2012), means that both the amplitude and detailed structure of CMB heat-flux variations remain poorly constrained (Koelemeijer, 2021). Assuming that the temperature on the core side of the CMB is approximately uniform, anomalously hot regions of the lowermost mantle reduce the heat flux out of the core due to the reduced temperature difference between the core and the mantle. In anomalously cold regions of the lowermost mantle, the heat flux out of the core is increased, producing regional heat-flux maxima rather than minima. These may be associated with cold downwellings or subducted slab remnants in the mantle and can impose persistent non-zonal thermal forcing on the geodynamo (Mound and

Davies, 2023; Dannberg et al., 2024).

Because the core is vigorously convecting, lateral temperature variations within the core are expected to be small, so the temperature at the top of the core is approximately uniform. Lateral temperature variations in the lowermost mantle can therefore produce spatial variations in heat flux across the CMB. These variations give rise to heterogeneous CMB heat-flux patterns that evolve on mantle overturn timescales, of order 10 Myr, which are long compared with characteristic core timescales. The imposed heat-flux pattern can therefore be treated as approximately static from the perspective of the core, providing a persistent thermal boundary condition on the geodynamo (Frasson et al., 2025; Dannberg et al., 2024).

Numerical simulations show that the boundary condition can affect the structure and evolution of the magnetic field. In some simulations, reduced heat flux beneath hot mantle anomalies can produce long-lived zones of quasi-stable fluid in the uppermost core, including inversion lenses, which may imprint persistent non-zonal structure on the geodynamo (Mound and Davies, 2023; Biggin et al., 2026). Lateral variations in CMB heat flux may also favour regional, rather than globally continuous, stratification, since a single stratified layer would be difficult to reconcile with radial upwellings at the CMB (Whaler, 1980; Huguet et al., 2018). Consistent with this, Mound et al. (2019) predict regional stratification arising from lateral heat-flux variations, while Naskar et al. (2026) show the generation of regional inversion lenses beneath thermal anomalies.

Together, these studies show that the structure of CMB heat flux can influence the morphology and evolution of the geomagnetic field. This influence may be reflected in the distribution of flux patches, which are localised regions of strong radial magnetic field at the CMB. However, their distribution does not provide unique evidence of a persistent mantle imprint, since long-term field models and archeomagnetic studies suggest that flux patches can be mobile and are not fully locked to preferred longitudes (Amit et al., 2011; Panovska et al., 2019). Seismic constraints on stratification at the top of the core are also debated (van Tent et al., 2020). What remains even less clear is whether the amplitude of this imposed heterogeneity can be recovered from the magnetic field itself. Larger heat-flux heterogeneity amplitudes may strengthen the mantle imprint on the core flow, with possible consequences for non-zonal field structure, persistent flux patches, secular variation, reversal behaviour and regional stratification. However, it is not clear whether these magnetic signatures are caused entirely by the mantle forcing, since similar behaviour may arise in models with different parameter choices or boundary conditions.

The extent of this non-uniqueness has not yet been quantified, making it difficult to isolate the influence of heat-flux heterogeneity amplitude alone.

To determine whether the amplitude of CMB heat-flux variations can be inferred from the radial magnetic field at the CMB, we apply a machine-learning approach to outputs from geodynamo simulations. Machine learning refers to a set of statistical techniques that use data to undertake a task, without being explicitly programmed to do so, allowing the user to extract knowledge and draw inferences from data (Jordan and Mitchell, 2015). Here, we use an autoencoder-based architecture. An autoencoder is a neural network that learns to compress input data into a lower-dimensional latent representation, or encoding, and then reconstruct the original input from this compressed representation (Bank et al., 2023). Here, latent refers to learned features that are not directly observed in the input data, but summarise the information needed to reconstruct it. The compression step forces the network to identify the most informative features needed to reproduce the input, which in image data may include structures such as shapes, edges, boundaries, textures, and broader spatial patterns. For classification, this latent representation can be coupled to a classifier by adding a classification layer or head, and training the network with a combination of reconstruction and classification losses. In this way, the latent space is encouraged to retain information that is useful both for reconstructing the magnetic-field input and for predicting the heat-flux heterogeneity amplitude class.

Autoencoder classifiers have been used in a diverse range of fields, including medicine (Ganotra et al., 2024; Mao et al., 2025), anomaly detection (Saeedi and Giusti, 2023; Su et al., 2015; Kozmin et al., 2025), and hyperspectral imaging (Xing et al., 2016). Alongside the development of these methods, advances have also been made in explainable AI (XAI), which aims to provide explanations for machine-learning outputs in forms that are understandable to humans (Doshi-Velez and Kim, 2017). One such technique is the use of saliency maps, which highlight the regions of an input, such as pixels in an image, that contribute most strongly to a model prediction (Linardatos et al., 2021). In the context of this study, saliency maps provide a way to investigate which parts of the CMB radial magnetic field are most important for distinguishing between different amplitudes of imposed heat-flux heterogeneity.

In this study, we present a convolutional autoencoder classifier, which builds on previous numerical studies of core–mantle thermal coupling by testing whether information about the amplitude of the imposed mantle forcing is encoded in the resulting magnetic-field morphology. Classifier

training is performed using a joint loss that combines reconstruction and classification objectives, and performance is evaluated using both reconstruction error and classification accuracy. The methodology is described in Section 6.2. In Section 6.3, we first apply the framework to the simulations of Mound and Davies (2023), using these experiments to validate the approach and to investigate the magnetic-field features contributing most strongly to each class of solutions. In Section 6.4, we extend the analysis to a broader set of Earth-like simulations, characterised by a mean magnetic-to-kinetic energy ratio $M_{\mathrm{ratio}} \geq 1$ and a mean magnetic Reynolds number $R_m$ in the range 400–2200. The application to the Earth's magnetic field is shown in Section 6.5, and the main discussion and conclusions are presented in Section 6.6.

## 6.2 Methodology

The classification methodology consists of a convolutional autoencoder with an integrated classifier head, implemented in PyTorch (Paszke et al., 2019). A schematic of the network architecture is shown in Figure 6.1. The autoencoder learns a compressed representation of the input magnetic field, while the classifier uses this latent representation to predict the class associated with the amplitude of the imposed CMB heat-flux heterogeneity, denoted by $Q^*$,

$$Q^* = \frac{Q_{max} - Q_{min}}{Q_{average}}, \tag{6.1}$$

where $Q_{\mathrm{max}}$, $Q_{\mathrm{min}}$, and $Q_{\mathrm{average}}$ are the maximum, minimum, and average heat fluxes per unit area at the CMB, respectively (Mound and Davies, 2023). Depending on the application, the classifier is trained either as a three-class classifier, distinguishing between simulations with different $Q^*$, or as a two-class classifier, distinguishing between the homogeneous case, where $Q^* = 0$, and all heterogeneous cases, where $Q^* \neq 0$.

For each input snapshot, the encoder first converts the input image of $B_r$ into a smaller set of learned feature maps. These feature maps form the latent representation, which is the compressed description of the magnetic-field image used by the network. The classifier uses this latent representation to predict the heat-flux class, while the decoder uses it to reconstruct the original magnetic-field image. The classifier is therefore trained to retain information that is useful both for describing the input field and for distinguishing between heat-flux classes.

The input to the network is a collection of two-dimensional images of the radial magnetic field, $B_r$, each taken from a numerical dynamo simulation evaluated at the CMB, and projected onto a regular latitude–longitude grid of size $40 \times 89$. Convolutional layers are designed for data with a known grid-like topology, such as images, where local filters are applied consistently across neighbouring pixels (Goodfellow et al., 2016). In this representation, the magnetic field is treated as a single-channel image. A channel is one layer of information stored at each grid point. Here, each grid point contains only one value, the radial magnetic field $B_r$, so the input has one channel. This differs from an image with several types of information at each pixel, such as separate red, green, and blue colour values. Since the longitudinal direction is periodic, the convolutional layers use periodic padding in longitude so that filters applied at one longitudinal edge can access values from the opposite edge. No periodic padding is applied in the latitudinal

direction.

The encoder consists of three convolutional layers, which progressively convert the input image into coarser feature maps while increasing the number of feature channels. These channels are feature maps, which show where each learned filter detects particular patterns in the input data. Rather than storing raw magnetic-field values, feature maps represent the classifier's internal extraction of underlying spatial features, such as edges, shapes, boundaries, or coherent magnetic-field structures. In this way, the encoder extracts spatial features from the magnetic-field image and maps them onto a lower-dimensional latent representation. The decoder consists of three transposed convolutional layers, which increase the spatial dimensions of the latent representation to recover the original input size. The encoder and decoder have separate learned weights, but skip connections pass feature maps directly from each encoder layer to the corresponding decoder layer, helping to reconstruct small-scale and high-resolution spatial features. The decoder therefore depends on both the latent representation and earlier encoder features.

The classifier head is kept deliberately simple so that the class prediction reflects the information contained in the encoded representation, rather than the capacity of an additional deep neural network. Global average pooling reduces each of the four final encoded feature maps to a single value by averaging over its spatial dimensions, giving a four-dimensional summary of the latent representation. This is then mapped to the three heat-flux classes using a single fully connected layer, which reduces the number of trainable parameters and helps limit over-fitting while keeping the classification pathway directly tied to the encoder output. This choice follows common convolutional neural network (CNN) classification practice, where global average pooling is used to summarise final feature maps before a simple linear classifier, reducing the dependence on larger fully connected classification networks (Zhou et al., 2016; He et al., 2016).

The network is trained by minimising a loss function that combines a reconstruction objective and a classification objective. The reconstruction loss, $L_r$, measures the Huber loss between the input magnetic field, $B_r$, and the reconstructed magnetic field. Huber loss behaves like squared-error loss for small errors, up to a threshold $\delta$, and like absolute-error loss for larger errors, reducing the influence of unusually large differences (Huber, 1964). The classification loss, $L_c$, measures the cross entropy loss between the predicted class and the true class label. Cross entropy loss compares the predicted class probabilities with the true class label, and penalises the classifier more strongly when it assigns a low probability to the correct class

(Goodfellow et al., 2016). The total loss, $L_{TOTAL}$, is defined as

$$L_{TOTAL} = \alpha L_r + (1 - \alpha) L_c, \tag{6.2}$$

where $\alpha$ controls the relative weighting of the two loss terms. We use $\alpha = 0.4$, which was found empirically to give the lowest total loss compared to other values of $\alpha$ varying between 0 and 1. This value may be specific to this dataset, and a different dataset may require a different value of $\alpha$. This choice places greater weight on the classification task, while still requiring the latent representation to retain enough information to reconstruct the input magnetic field.

Before training, the dataset was divided into a training set and a validation set. The training set was used to update the network weights, while the validation set was withheld from weight updates and used to monitor performance during training. The classifier was trained using the AdamW optimiser (Loshchilov and Hutter, 2017) for 100 iterations. After each iteration, the reconstruction and classification losses were evaluated on the validation set, and the checkpoint with the lowest total validation loss was retained. The validation set therefore contains magnetic-field snapshots that were not used to update the network weights, although it was used to select the preferred checkpoint. The reported validation errors provide an indication of performance on withheld snapshots, but may be influenced by this classifier-selection step.

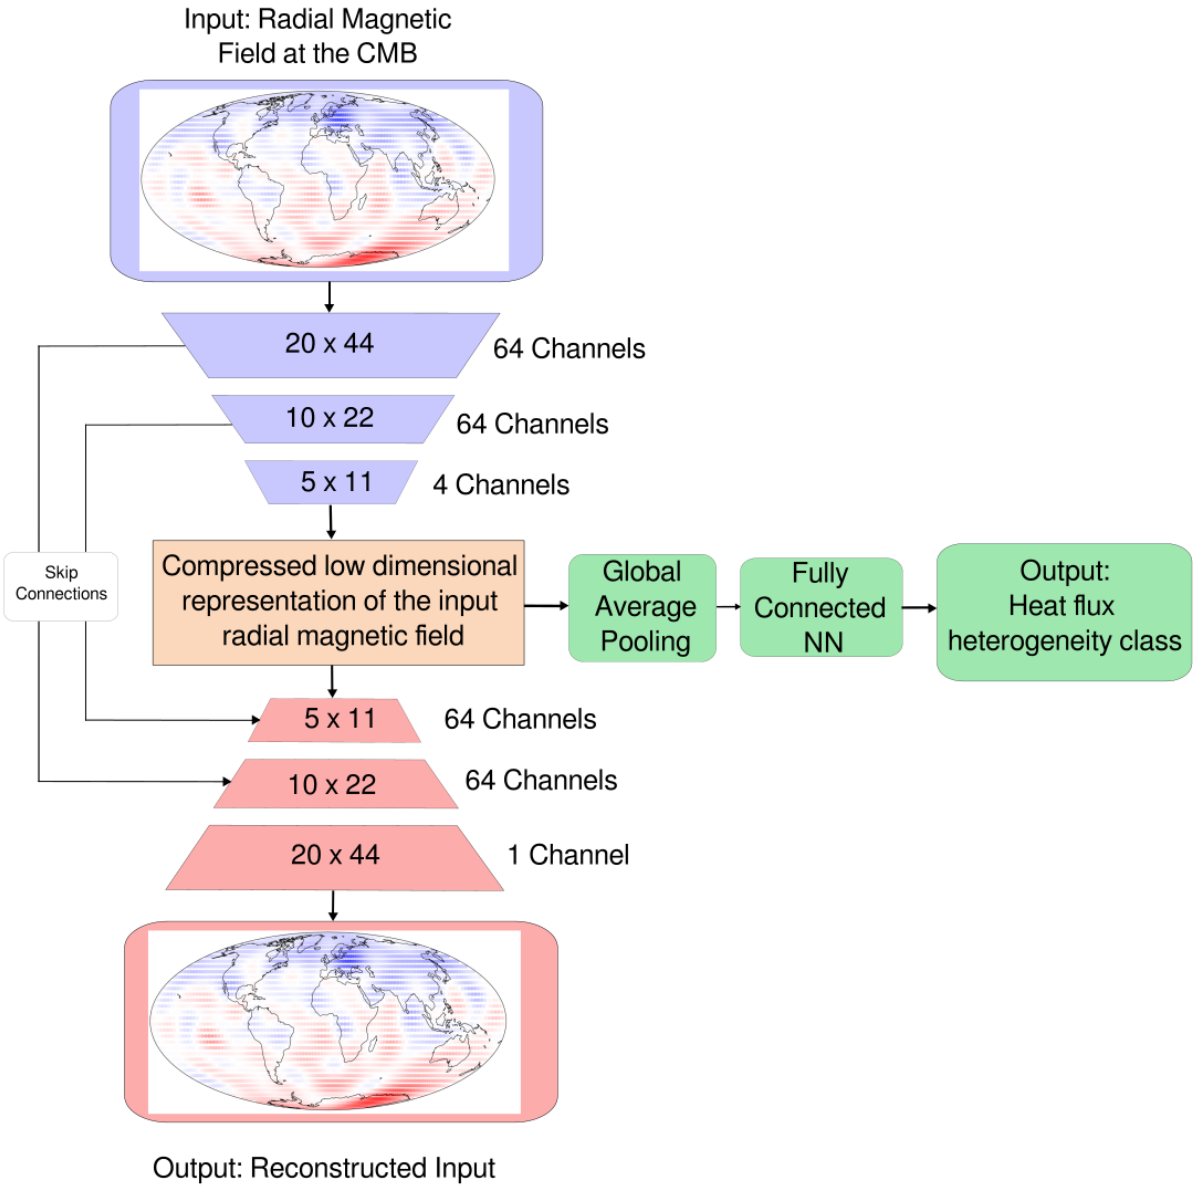



Figure 6.1: Schematic of the convolutional autoencoder classifier used in this study. Images of the radial magnetic field at the CMB are encoded into a low-dimensional latent representation, which is passed to the classifier to predict the heat-flux heterogeneity class. The decoder reconstructs the original radial magnetic field from the latent representation. The number of channels in each layer was chosen by considering the compression ratio between the input and latent images, which is described in Section 6.3.2.

## 6.3 Classification using simulations in Mound and Davies (2023)

### 6.3.1 Training Data

As an initial proof of concept, we first apply the classification framework to a limited subset of numerical dynamo simulations. This allows the method to be tested on a controlled set of cases before being extended to a broader parameter range. We use the simulations presented in Mound and Davies (2023), which differ in the amplitude of the thermal heterogeneity imposed at the CMB, quantified by $Q^*$, and in the Rayleigh number, $Ra$, which measures the strength of thermal driving relative to viscous and thermal diffusion. The dataset uses two values of the Rayleigh number, $Ra = 200000000$ and $Ra = 600000000$, and three values of imposed heat-flux heterogeneity, $Q^* = 0$, 2.46, and 5.35. Several other control parameters are fixed across this set of simulations. The Ekman number, $E_k$, which measures the ratio of viscous to Coriolis forces, is set to $E_k = 10^{-5}$. The magnetic Prandtl number, $P_m$, defined as the ratio of kinematic viscosity to magnetic diffusivity, is set to $P_m = 1$, and the Prandtl number, $P_r$, which gives the ratio of momentum diffusivity to thermal diffusivity, is fixed at $P_r = 0.2$. Although a simpler test would be to fix $Ra$ and classify only between simulations with different $Q^*$, we include both values of $Ra$ so that the proof-of-concept test is not restricted to a single convective regime. The heterogeneous cases use the same tomography-derived spatial pattern based on Masters et al. (1996), with its amplitude controlled by $Q^*$, while the $Q^* = 0$ cases have a homogeneous outer-boundary heat flux.

A summary of the simulations used in this study is given in Table 6.1. The class-averaged radial magnetic field at the CMB in the equatorial region for each value of $Q^*$, for all values of $Ra$, is shown in Figure 6.2, and the corresponding class-averaged radial magnetic field in the polar region is shown in Figure 6.3. In these class-averaged fields, increasing $Q^*$ is associated with an increase in the intensity of the radial magnetic field. In the equatorial region, shown in Figure 6.2, this is expressed as stronger magnetic-field features beneath the Pacific, the Indian Ocean, and South America. In the polar region, shown in Figure 6.3, increasing $Q^*$ is associated with the development of high-latitude flux patches and polar minima. These systematic differences in the radial magnetic field provide the information available to the classifier when distinguishing between simulations with different values of $Q^*$.

The training data are constructed from the six resulting $Q^*$–$Ra$ combinations. As the simulations differ in non-dimensional duration and in the number of saved state files, 2000 snapshots were selected at random from each simulation. The snapshots are therefore not evenly spaced in non-dimensional time and may retain some temporal correlation, but random sampling avoids restricting the dataset to a particular part of each simulation. The final dataset is balanced, with 2000 snapshots per $Q^*$–$Ra$ combination and 12,000 snapshots in total.

| **Simulation Name** | $Q^*$ | $Ra$ | **Number of Snapshots** |
|---|---|---|---|
| LEDCJ001 | 0 | 200000000 | 2000 |
| LEDCJ002 | 2.46 | 200000000 | 2000 |
| LEDCJ003 | 5.35 | 200000000 | 2000 |
| LEDCJ004 | 0 | 600000000 | 2000 |
| LEDCJ005 | 2.46 | 600000000 | 2000 |
| LEDCJ006 | 5.35 | 600000000 | 2000 |

Table 6.1: Summary of the simulations used for training (Mound and Davies, 2023).

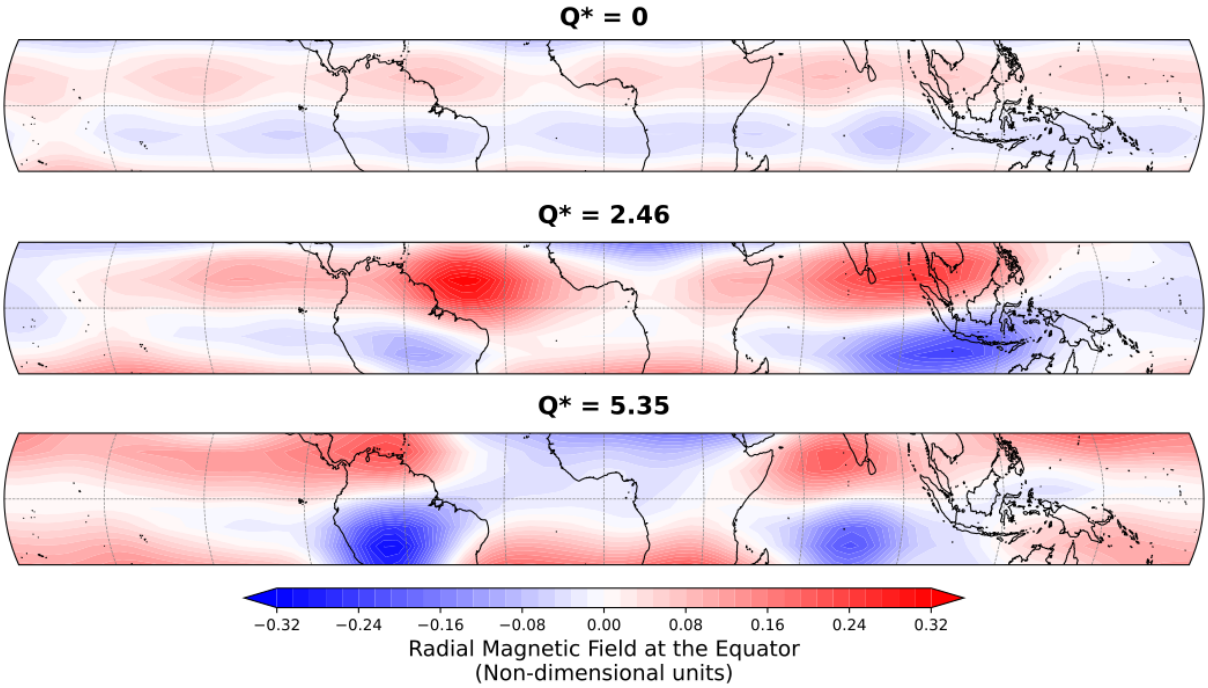



Figure 6.2: Class-averaged radial magnetic field in the equatorial region, truncated at spherical harmonic degree 12. The top panel shows $Q^* = 0$, the middle panel shows $Q^* = 2.46$, and the bottom panel shows $Q^* = 5.35$.

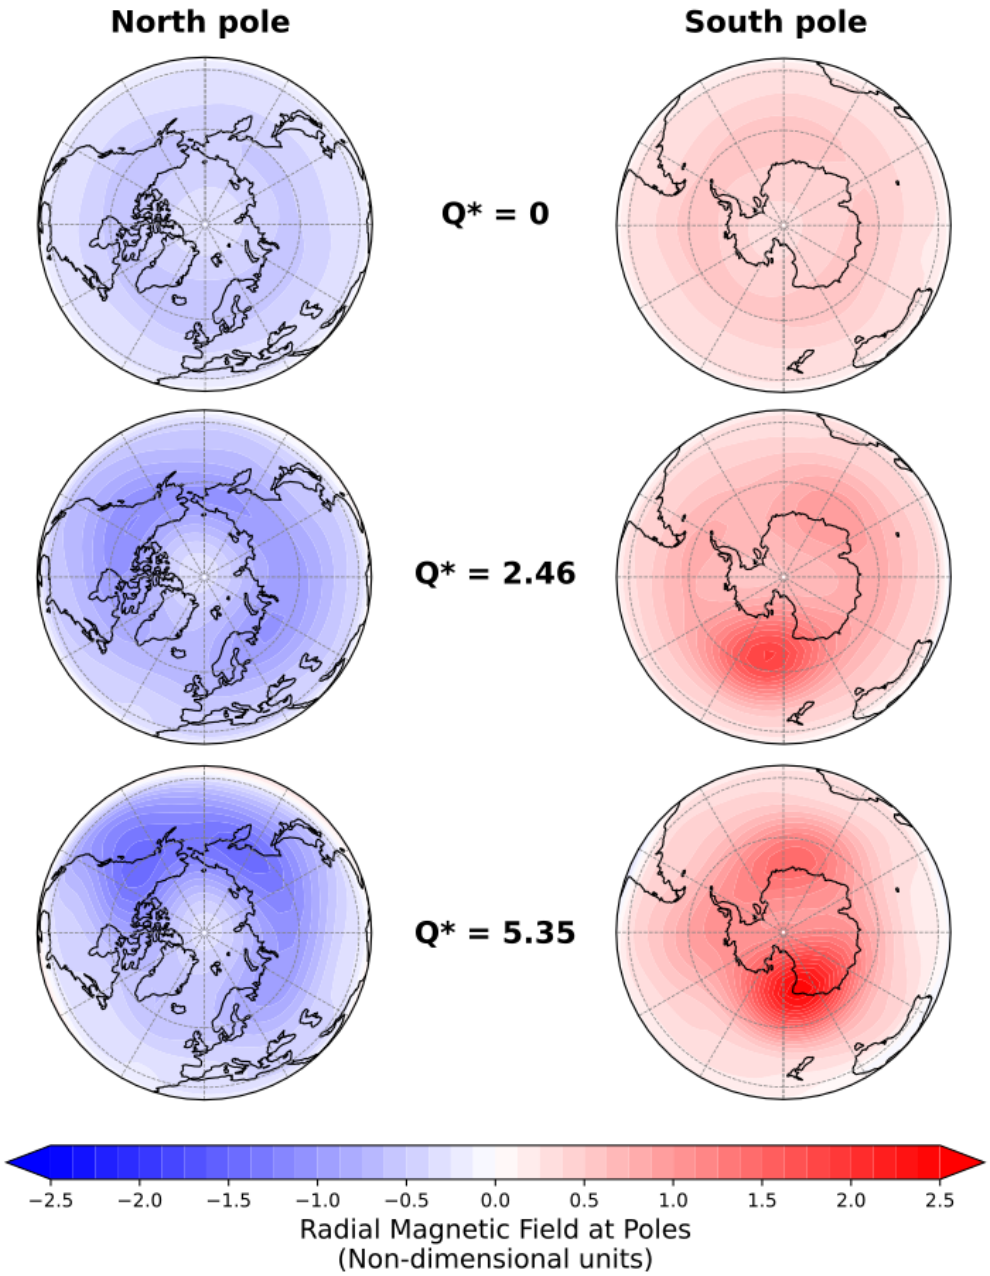



Figure 6.3: Class-averaged radial magnetic field at the poles, truncated at spherical harmonic degree 12. The top panel shows $Q^* = 0$, the middle panel shows $Q^* = 2.46$, and the bottom panel shows $Q^* = 5.35$.

### 6.3.2 Hyperparameter Tuning

The effects of input scaling, maximum spherical harmonic degree, compression ratio and temporal partitioning were assessed through a series of tuning experiments. Each was varied individually while the remaining choices were held fixed. The value giving the lowest validation error was then used in subsequent experiments. This approach allows the effect of each choice to be compared directly, but assumes that the choices can be assessed independently and does not account for possible interactions between them.

#### Scaling

The radial magnetic field input to the network was re-scaled so that images from different simulations could be compared on a similar scale (Sola and Sevilla, 1997). This is important because the framework is intended to be applied across simulations with different magnetic-field magnitudes, and may eventually be extended to observationally derived models of the Earth's magnetic field. Several scaling strategies were tested, including per-snapshot approaches that remove relative amplitude information and dataset-level approaches that retain amplitude differences.

In the first approach, each individual snapshot was normalised using the mean, $\mu$, and standard deviation, $\sigma$, of the radial magnetic field at that snapshot, such that

$$B_{\text{Normalised}} = \frac{B_r - \mu}{8\sigma}, \tag{6.3}$$

where $B_r$ is the radial magnetic field. The factor of 8 was chosen so that $B_{\text{Normalised}}$ lay approximately between $-1$ and 1. This scaling removes information associated with the absolute amplitude of the magnetic field at each snapshot.

In the second approach, each individual snapshot was normalised using the mean and standard deviation calculated across all snapshots, denoted by $\overline{\mu}$ and $\overline{\sigma}$, rather than values calculated separately for each snapshot. In this case, the radial magnetic field was scaled according to

$$B_{\text{Normalised}} = \frac{B_r - \overline{\mu}}{8\overline{\sigma}}. \tag{6.4}$$

Unlike the scaling in Eq. 6.3, this approach retains information associated with amplitude

differences between simulations.

The third and fourth approaches were based on scaling by the average magnetic energy at the CMB. This was calculated from the spatial power spectrum of the magnetic field which is given at the CMB by

$$E_{\mathrm{CMB}} = \sum_{n=1}^{n_{\max}} (n+1) \left(\frac{a}{r}\right)^{2n+4} \sum_{m=0}^{n} \left[(g_n^m)^2 + (h_n^m)^2\right], \tag{6.5}$$

where $g_n^m$ and $h_n^m$ are the spherical harmonic coefficients, and $a$ is the radius of the CMB (Lowes, 1974). The radial magnetic field was then scaled as

$$B_{\mathrm{Normalised}} = \frac{B_r}{\sqrt{E_{\mathrm{CMB}}}}. \tag{6.6}$$

Two variants of this energy-based scaling were tested. In the first, each individual snapshot was scaled using the energy at that snapshot. In the second, all snapshots were scaled using a single average energy calculated across all snapshots.

For each scaling strategy, classifiers were trained using radial magnetic field images truncated at different maximum spherical harmonic degrees, with the results shown in Figure 6.4. This analysis was carried out for both the three-class classification problem, distinguishing between $Q^* = 0$, $Q^* = 2.46$, and $Q^* = 5.35$, and the two-class problem distinguishing homogeneous from non-homogeneous cases. All scaling approaches showed similar behaviour, with classification errors for both the training set and the validation set increasing substantially for radial magnetic field images truncated at spherical harmonic degrees below 4. This suggests that smaller-scale magnetic-field structure, with $n_{max} \geq 4$, is required to distinguish thermal heterogeneity at the CMB. In addition, approaches using a single average scaling value performed better than those using values calculated separately for each snapshot, indicating that amplitude information contributes to the classification. The training data are truncated at spherical harmonic degree $n_{\max} = 12$, which is approximately the resolution of the core field. This choice is motivated by the eventual aim of applying the method to Earth observations, for which the internal field is most reliably resolved at low spherical harmonic degrees. The best performance at this spherical harmonic degree was obtained using the average energy scaling, Equation 6.6, which will be used for the remainder of this study.

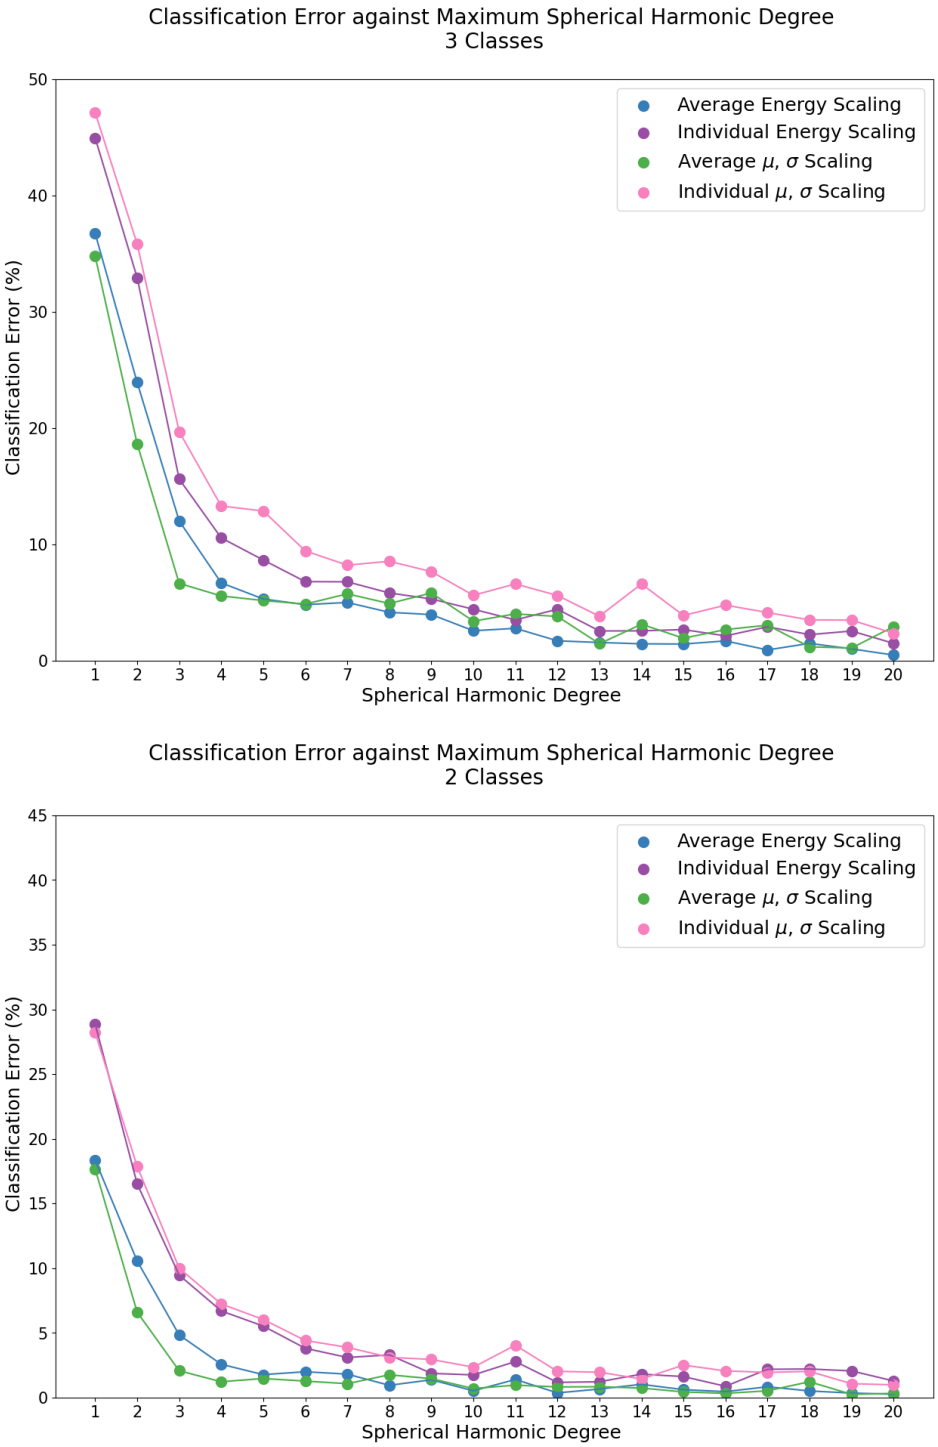



Figure 6.4: Classification error (%) for the validation dataset against spherical harmonic degree, for the three-class approach (top) and the two-class approach (bottom).

### Compression Ratio

A key choice in the autoencoder-classifier classifier is the size of the compressed representation used by the classifier. The autoencoder reduces each input magnetic-field image to a smaller set of values, known as the latent representation, from which the original image is reconstructed. In this study, the amount of compression is described using the compression ratio,

$$\mathrm{CR} = \frac{N_{\mathrm{input}}}{N_{\mathrm{latent}}}, \tag{6.7}$$

where $N_{\mathrm{input}}$ is the number of values in the input magnetic-field image, and $N_{\mathrm{latent}}$ is the number of values in the compressed latent representation.

The compression ratio determines how much information the classifier is allowed to retain before making the classification. A small compression ratio means that the latent representation is relatively large, so the classifier retains more detail from the input image. A large compression ratio means that the latent representation is smaller, forcing the classifier to describe the magnetic field using fewer values. This can help remove redundant information, but if the compression is too strong, features that are important for distinguishing between heat-flux heterogeneity classes may be lost.

In this architecture, the classifier is applied to the latent representation after global average pooling. The class prediction is therefore made from the compressed description of the radial magnetic field. The decoder also receives skip connections from earlier encoder layers, which help improve the reconstruction of the input image. However, these skip connections are used only by the decoder and are not passed to the classifier. The compression ratio reported here therefore describes the compactness of the representation used for classification, rather than the total information available to the reconstruction pathway.

To assess the effect of the compression ratio, values between CR = 1 and CR = 54 were tested. For each value, the classifier was trained using five different random initialisations of the network weights, and the resulting classification errors were recorded. The results are shown in Figure 6.5, where the mean classification error is shown in blue and the individual results from each random initialisation are shown in black. Within the tested range, CR = 13 gave the best balance between compression and classification accuracy. This value lies close to the knee of the trade-off curve, where the input representation has been substantially reduced without

producing a large increase in classification error, and also gives the lowest median classification error. Although low compression ratios, including CR = 1, gave similar classification errors, they retain a representation close to the full input image. This may allow the classifier to rely on details specific to the training snapshots, rather than on magnetic-field structures that generalise to unseen data. We therefore use CR = 13 for the results presented in this study, which gives comparable or better classification performance while forcing the classifier to classify from a reduced latent representation, encouraging it to retain more robust class-related features.

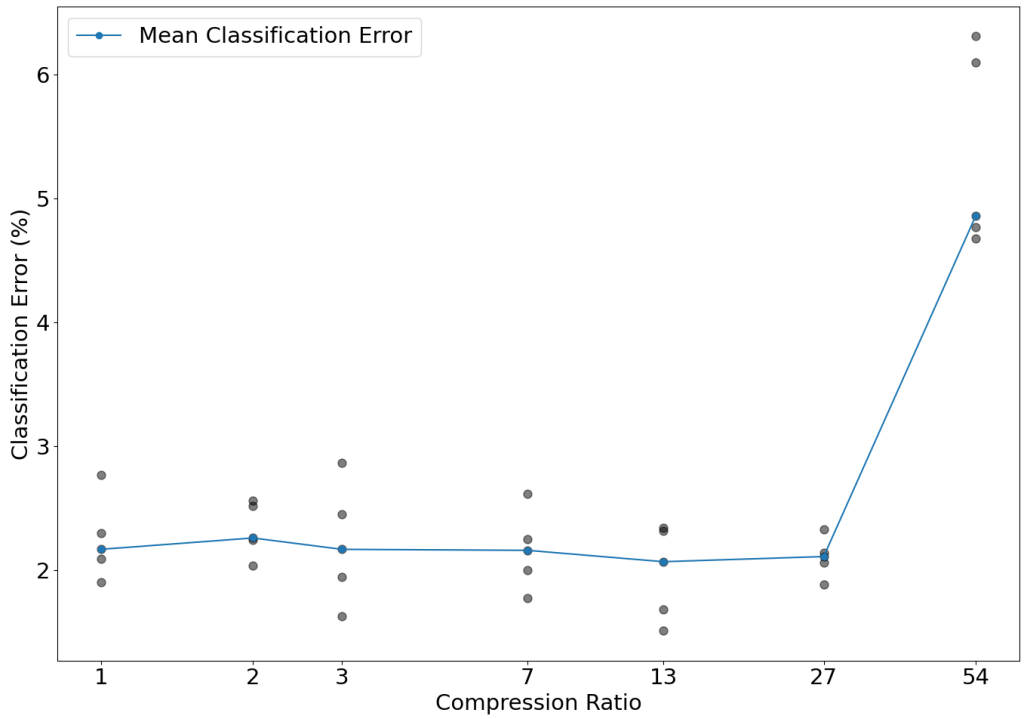



Figure 6.5: Classification percentage error against compression ratio. The mean classification error is shown in blue, and the individual results from each random initialisation are shown in black.

**Effect of Temporal Partitioning on Classification Error**

To assess the impact of simulation evolution on classification error, three strategies were considered for splitting the dataset into training and validation sets, where the training set is used to update the network weights and the validation set is withheld from training to assess generalisation to new magnetic-field snapshots. In the three-class configuration, the first approach involved randomly shuffling the data prior to partitioning into training and validation subsets. This approach yielded the lowest error rates, with classification errors of 1.5% and 4.9% for the training and validation sets, respectively. That this approach yielded the lowest error rates is unsurprising given that both subsets were drawn from the same underlying temporal distribution. In the second approach, the classifier was trained on samples from the first half of the simulations and evaluated on the second half. This resulted in poorer performance compared to the shuffled case for both the training and validation sets, with classification errors of 5.2% and 17.1%. This suggests the presence of temporal features in the latter half of the simulations that were not encountered during training, leading to reduced generalisation performance when the classifier operated outside its learned domain. The third approach reversed this configuration, training on the second half of the simulations and validating on the first. The resulting performance was broadly similar, with classification errors of 4.6% and 15.2%, marginally better than that obtained using the second approach.

A similar pattern is observed for the two-class configuration, with the shuffled results being 0.8% and 3.7% for the training and validation sets respectively, the first non-shuffled result being 3.8% and 9.9% for the training and validation sets respectively, and the second non-shuffled result being 1.6% and 8.4% for the training and validation sets respectively. This smaller value suggests that temporal differences primarily affect the classifier's ability to distinguish between varying degrees of heterogeneity, rather than between homogeneous and heterogeneous states outright. Notably, even under the worst-performing configuration the classifier achieves an accuracy of approximately 83%, indicating robust performance despite temporal effects. Based on these results, the shuffled partition was adopted for the subsequent experiments because it provided a consistent basis for comparing classifier configurations and gave the lowest validation error. The temporally partitioned results provide a more conservative indication of performance when the classifier is applied outside the temporal interval represented in the training data.

### 6.3.3 Sensitivity to regions of the CMB

To investigate which regions of the CMB contain the most useful information for classification, we test the spatial sensitivity of the classifier in two ways. This helps determine whether the classifier is using magnetic-field structures that are localised near regions of imposed heat-flux heterogeneity, or whether successful classification depends on broader features of the radial magnetic field. The first approach is to restrict the training and validation data to selected regions of the CMB and re-train the classifier using only these regions, allowing classification performance to be compared when different parts of the radial magnetic field are included or removed. The second approach uses saliency maps to identify which regions of the input field contribute most strongly to the predictions made by the trained classifier.

#### Restricting Input Data

We first consider the regional training tests. The classifier was retrained using the three-class approach, with restricted regions of the CMB used as input. The regions tested are shown by the coloured overlays in Figure 6.6. For each case, the trained classifier was then evaluated on validation data, selected using the shuffled splitting strategy. Regions for which the validation classification error was below 5% are shown in green, regions with errors between 5% and 10% are shown in orange, and regions with errors above 10% are shown in red.

The largest validation classification error was obtained using the region within 14 degrees of the poles, shown in Figure 6.6C, giving an error of 24.3%. Although this is the poorest-performing case, the error is still lower than would be expected from random classification, which for a 3-class test would be 66.7%, indicating that these regions do contain some information related to the imposed heat-flux heterogeneity. However, the polar regions alone do not appear to contain enough information to classify the simulations accurately.

Interestingly, the lowest validation classification error was not obtained when the full CMB was used, but when the regions within 14 degrees of the poles were excluded, shown in Figure 6.6F. The difference between this low- or mid-latitude case and the full-CMB case was small, with a maximum difference in validation classification error of 0.2%. This suggests that the classifier can still discriminate between the different heat flux heterogeneities even in the absence of the poles, and that these instead introduce variation that is not strongly related to the heat-flux heterogeneity amplitude. Removing these regions may therefore allow the classifier to focus

more strongly on the parts of the magnetic field that vary systematically between classes.

Overall, these results indicate that the equatorial and mid-latitude regions contain most of the information used by the classifier. classifiers trained using these regions gave the lowest validation classification errors, suggesting that the magnetic-field structures most strongly associated with the imposed heat-flux heterogeneity amplitude are concentrated away from the polar regions. This suggests that, for this set of simulations, the classifier is primarily using magnetic-field structure in the equatorial and mid-latitude CMB to distinguish between different values of $Q^*$.

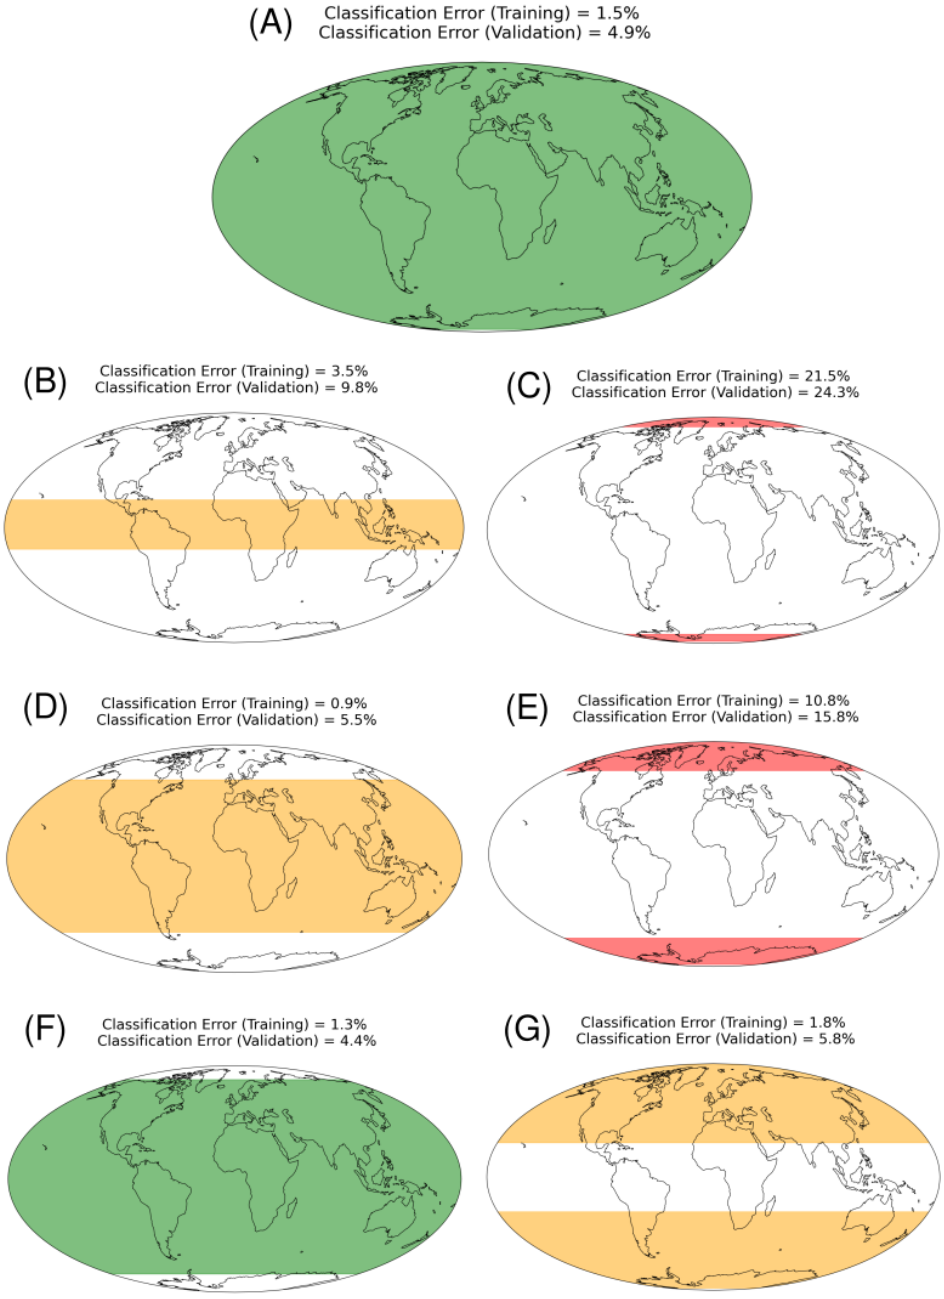



Figure 6.6: Classification Error when different regions (coloured overlay) were used for training data. Regions for which the validation classification error was below 5% are shown in green, regions with errors between 5% and 10% are shown in orange, and regions with errors above 10% are shown in red.

### 6.3.4 Saliency Maps

To further investigate which regions of the CMB were most important for classification, we applied Gradient-weighted Class Activation Mapping (Grad-CAM). Grad-CAM is a visualisation method that highlights the regions of an input image that have the greatest influence on a classifier prediction (Selvaraju et al., 2017). It uses the gradients of the predicted class score with respect to the final convolutional feature maps to estimate which regions most strongly affect the classification. These gradients are then used to weight the feature maps, producing a heat map of the regions most important for the classifier prediction. In this study, it is used to identify which parts of the radial magnetic field, $B_r$, the classifier relies on most strongly when assigning a snapshot to a particular class of $Q^*$. The maps were computed with respect to the class predicted by the classifier, so that they show the regions that contributed to the decision the classifier actually made.

Saliency maps were generated for all samples in the validation dataset. An example for a single sample is shown in Figure 6.7, with the radial magnetic field plotted underneath. Darker regions in the saliency map indicate areas of the input that had a stronger influence on the classification decision. To identify the regions that were important across the dataset as a whole, the saliency maps were then averaged over all samples with the same true and predicted class. The aggregate maps for the three-class classification are shown in Figure 6.8, using the same layout as a confusion matrix (Dobson, 2024). Maps on the diagonal, highlighted with a green background, correspond to correctly classified samples, while the off-diagonal maps correspond to cases in which the classifier made an incorrect prediction.

For the three-class classification, correctly classified homogeneous cases where both the predicted and true class correspond to $Q^* = 0$, show the strongest saliency around the poles and high latitudes, particularly in the northern hemisphere. For correctly classified intermediate-heterogeneity cases, corresponding to $Q^* = 2.46$, the strongest contributions are found in the regions under the southern Pacific Ocean, with some weaker contributions beneath South America and the Indian Ocean. Comparing these regions with the imposed heat-flux pattern shown in Figure 6.9 suggests that the classifier is sensitive to magnetic-field structure close to the Pacific LLVP, or to the lateral gradients associated with its boundaries. For correctly classified high-heterogeneity cases, corresponding to $Q^* = 5.35$, the strongest saliency is concentrated mainly in the region of the Pacific LLVP. The southern Pacific is a significant feature in both

heterogeneous classes, and also corresponds to the region of largest imposed heat-flux anomaly, suggesting that magnetic-field structure in this region is particularly sensitive to the imposed boundary heterogeneity and helps the classifier distinguish heterogeneous from homogeneous heat-flux cases.

The saliency maps for incorrectly classified samples provide information about why the classifier may have made these errors. In these cases, the maps tend to highlight regions associated with the predicted class rather than the true class. For example, when the classifier predicts the intermediate-heterogeneity class, $Q^* = 2.46$, but the true class is the homogeneous case, $Q^* = 0$, the saliency is concentrated in equatorial regions associated with the LLVPs. Similarly, when the classifier predicts $Q^* = 2.46$ but the true class is $Q^* = 5.35$, the saliency pattern resembles that of correctly classified $Q^* = 2.46$ samples. This suggests that misclassifications occur when features resembling those of another class dominate the classifier decision, causing physically different magnetic fields to appear similar to the classifier.

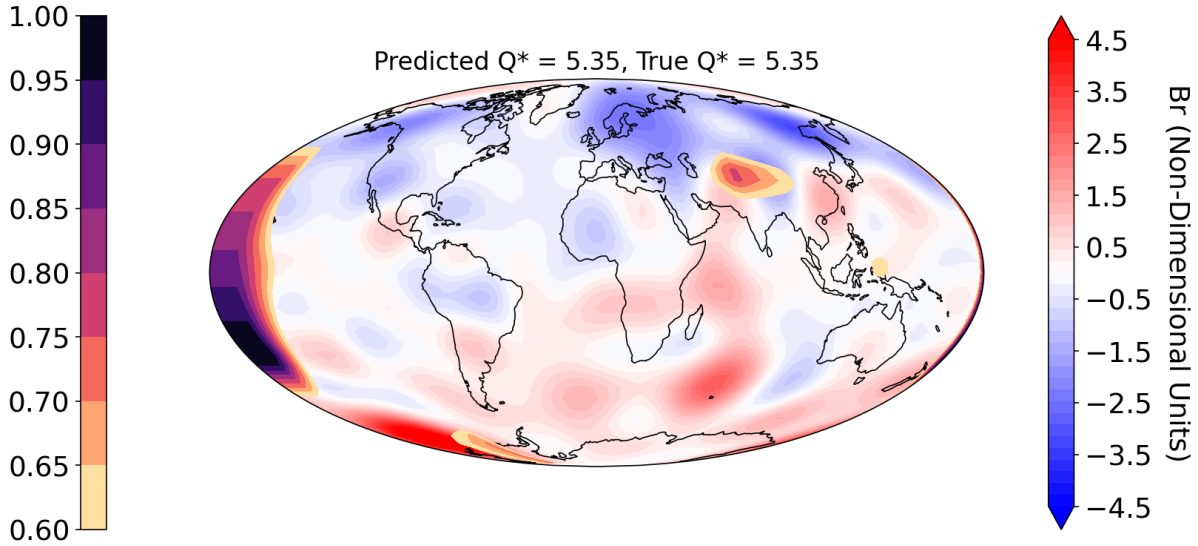



Figure 6.7: Example of an individual saliency map, with the radial magnetic field plotted on the same grid but using a different colour map. Darker regions (using the colourbar on the left) indicate areas of the input that contributed more strongly to the classifier prediction. The saliency maps are normalised and therefore dimensionless, with larger values indicating regions that contributed more strongly to the classifier prediction relative to the rest of the map.

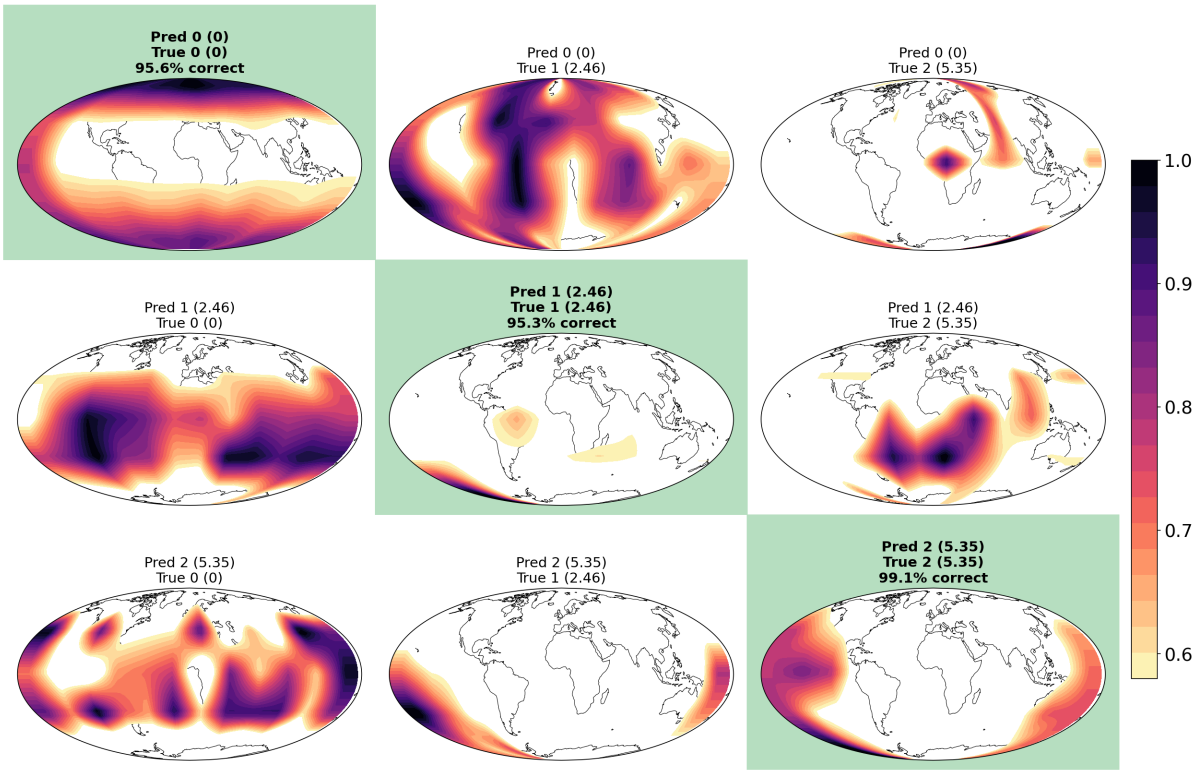



Figure 6.8: Aggregate saliency maps for the three-class classification approach, shown in a confusion-matrix layout. Maps on the diagonal, with a green underlay, correspond to correctly classified samples, while off-diagonal maps correspond to misclassified samples. Darker regions indicate areas that contributed more strongly to the classifier prediction.

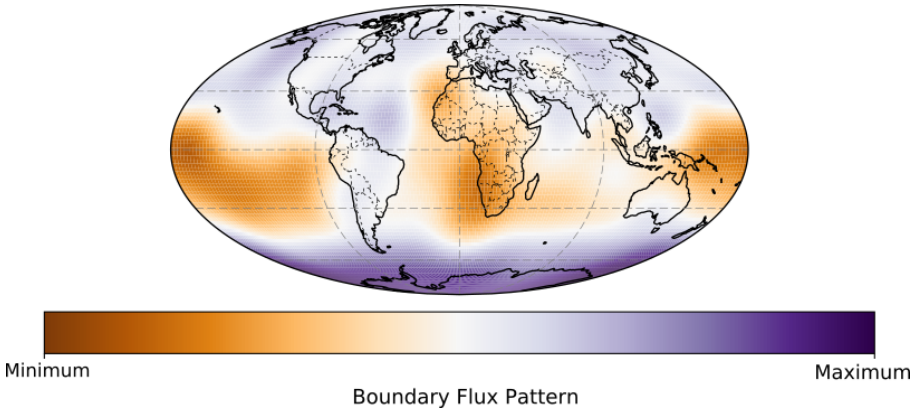



Figure 6.9: Map showing the heat flux pattern of the LLVPs applied to the training simulations.

A similar pattern is seen for the two-class classification, shown in Figure 6.10. Correctly classified homogeneous cases show the strongest saliency in the polar regions and the Pacific. For correctly classified heterogeneous cases, the strongest saliency is found at high southern latitudes, close to the Pacific LLVP. Comparison with individual saliency maps, such as those shown in Figure 6.11, suggests that these regions often coincide with high-flux patches near the pole. Since this area is highlighted for heterogeneous cases in both the two-class and three-class classifications, magnetic-field structure in the southern Pacific may be one of the features used by the classifier to separate homogeneous and heterogeneous simulations.

Taken together, the regional training tests and saliency maps suggest that the classifier uses spatially distributed information from the radial magnetic field. Both approaches were evaluated using the same validation set, but they probe the classifier in different ways. The regional tests measure how classification performance changes when selected parts of the CMB are included or removed, whereas the saliency maps identify the regions that contribute most strongly to the trained classifier's predictions. Both approaches indicate that no single region is sufficient by itself. Instead, successful classification requires information from a range of regions, particularly at equatorial and mid-latitudes near the Pacific and African LLVPs. This is consistent with the compression-ratio tests in Section 6.3.2, where the best performance was obtained using a relatively large latent representation. If the classification depended on a single localised feature, a smaller latent space might be expected to perform well. The larger latent space instead suggests that the classifier retains information from several spatial features of the magnetic field.

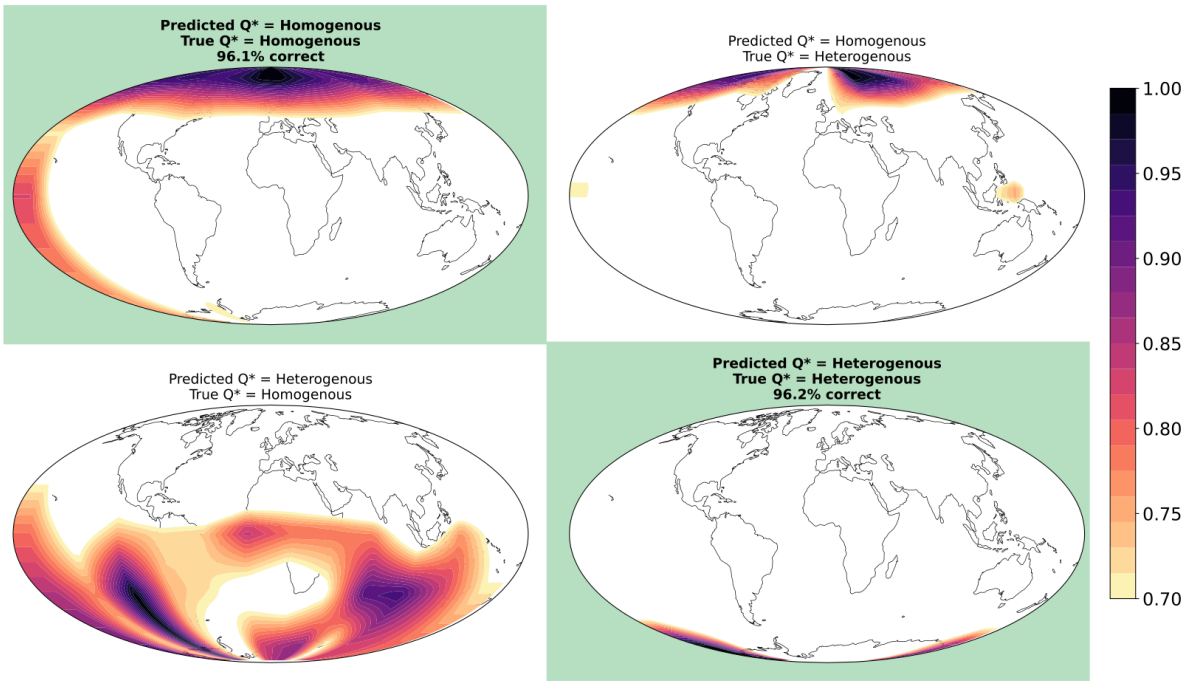



Figure 6.10: Aggregate saliency maps for the two-class classification approach, shown in a confusion-matrix layout. Maps on the diagonal, with a green underlay, correspond to correctly classified samples, while off-diagonal maps correspond to misclassified samples. Darker regions indicate areas that contributed more strongly to the classifier prediction.

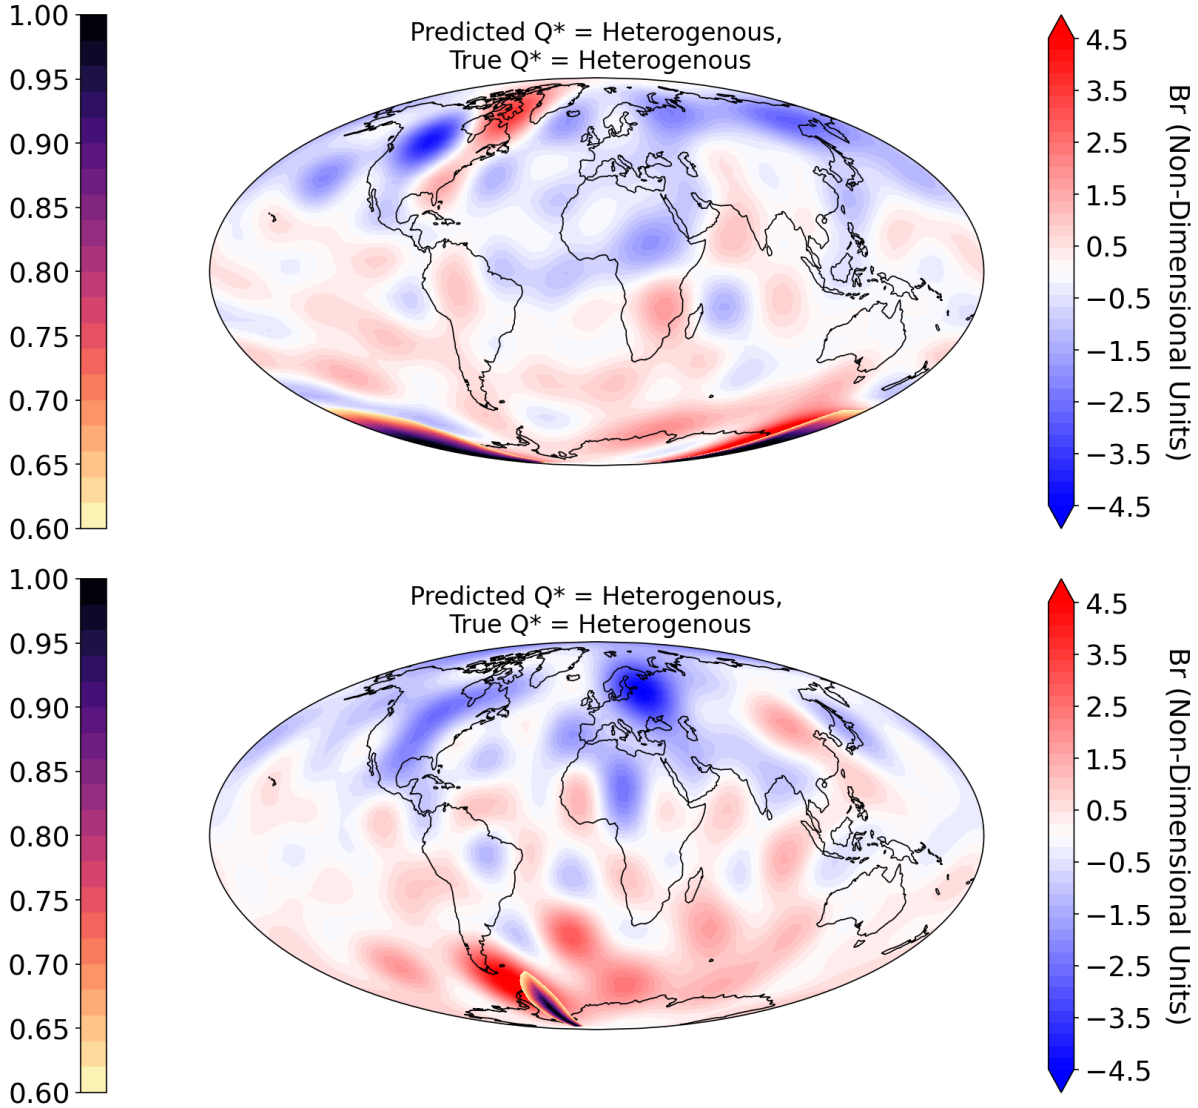



Figure 6.11: Examples of individual saliency maps for a $Q^* \neq 0$ 2-class case, with the radial magnetic field plotted alongside. The saliency maps consistently highlight flux patches near the southern pole, such as in these examples, as areas that contribute more strongly to the classifier prediction.

## 6.4 Classification using a wider set of simulations

Having tested our methodology on a limited set of simulations, we then applied the method to a wider set of simulations, taken from the set published in Rogers et al. (2026). We did not run any simulations specifically for this study. We selected simulations with $400 \leq R_m \leq 2200$, since this overlaps estimates for Earth's core, where $R_m$ is expected to be $O(10^3)$ and has been estimated to lie between roughly 400 and 4000 (Aubert et al., 2017; Nakagawa and Davies, 2022). We also required a magnetic to kinetic energy ratio $M_{\text{ratio}} \geq 1$, which is lower than Earth's expected value but selects simulations in which magnetic energy exceeds kinetic energy, giving access to the strong-field regime associated with QG-MAC-like dynamics (Aubert et al., 2017; Nakagawa and Davies, 2022). These criteria therefore select simulations that approach Earth-like magnetic induction and force-balance behaviour, while remaining computationally accessible. Ideally, one would use a narrow suite of simulations with Earth-like parameters that differ only in $Q^*$, so that changes in the magnetic field could be attributed directly to the imposed heat-flux heterogeneity. However, the appropriate parameter values for the Earth are uncertain, and numerical simulations cannot reach all Earth-like limits. A wider simulation set is therefore useful because it tests whether the classifier can identify heat-flux heterogeneity despite uncertainty in the underlying dynamo regime.

All other parameters were allowed to vary, and the $Q^*$ values ranged from 0.13–10.72. A summary of the simulations used in the training data is given in Table 6.2, while information on heterogeneity patterns and references is given in the Appendix. These simulations were chosen as they were all published, or had formulations that had been published, and allowed for the number of snapshots in each class to be similar. We applied only the two-class approach to these simulations, due to limitations in the number of simulations available for each individual value of $Q^* \neq 0$. Again, as the simulations differ in non-dimensional duration and therefore in the number of saved state files, the snapshots were selected at random from each simulation, giving 7954 homogeneous snapshots and 7982 heterogeneous snapshots. The resulting dataset was therefore approximately balanced between the two classes before the train–validation split. However, the heterogeneous class was not uniformly distributed in $Q^*$, because many of the available heterogeneous simulations had low imposed heat-flux heterogeneity amplitudes. The Class-averaged radial magnetic field in the equatorial region is shown for each value of $Q^*$ in Figure 6.12, with the corresponding polar views shown in Figure 6.13. The wider simulation set

contains a larger proportion of cases of reversals, together with greater morphological variability in $B_r$. As a result, these averages have lower field intensity and differ in structure from those shown in Figures 6.2 and 6.3. As $Q^*$ is increased from the homogeneous case, more concentrated flux patches develop, particularly near the South Pole. All radial magnetic-field maps were normalised using the mean magnetic energy at the CMB across the full simulation set. The autoencoder used a compression ratio of 13, as in Section 6.3.

Table 6.2: Summary of the simulations used for training data.

| **Simulation Name** | $Ra$ | $Ek$ | $Rm$ | $M_{\text{ratio}}$ | $Q^*$ | **Number of Snapshots** |
|---|---|---|---|---|---|---|
| LEDCJ001 | 200000000 | 0.000020 | 888 | 2.2 | 0.0 | 194 |
| LEDSM002 | 1200000 | 0.00040 | 1234 | 12 | 0.0 | 194 |
| LEDT001 | 120000 | 0.0020 | 1127 | 2.6 | 0.0 | 194 |
| LEDT002 | 480000 | 0.0010 | 1196 | 3.2 | 0.0 | 194 |
| LEDT002_SM | 480000 | 0.0010 | 1192 | 3.2 | 0.0 | 194 |
| LEDT003 | 1920000 | 0.00050 | 1379 | 3.2 | 0.0 | 194 |
| LEDT003_SM | 1920000 | 0.00050 | 1354 | 3.6 | 0.0 | 194 |
| LEDT004 | 7680000 | 0.00025 | 1543 | 3.5 | 0.0 | 194 |
| LEDT008 | 360000 | 0.0010 | 936 | 6.5 | 0.0 | 194 |
| LEDT008_SM | 360000 | 0.0010 | 969 | 5.1 | 0.0 | 194 |
| LEDT009 | 1100000 | 0.00050 | 962 | 6.7 | 0.0 | 194 |
| LEDT009_SM | 1100000 | 0.00050 | 960 | 6.6 | 0.0 | 194 |
| LEDT010 | 3360000 | 0.00025 | 966 | 7.3 | 0.0 | 194 |
| LEDT010_SM | 3360000 | 0.00025 | 955 | 7.6 | 0.0 | 194 |
| LEDT013 | 4800000 | 0.00020 | 959 | 7.9 | 0.0 | 194 |
| LEDT013_SM | 4800000 | 0.00020 | 952 | 8.0 | 0.0 | 194 |
| LEDT014_SM | 10080000 | 0.00013 | 925 | 9.2 | 0.0 | 194 |
| LEDT015 | 14500000 | 0.00010 | 935 | 8.9 | 0.0 | 194 |
| LEDT015_SM | 14500000 | 0.00010 | 918 | 10 | 0.0 | 194 |
| LEDT016 | 30720000 | 0.000063 | 893 | 12 | 0.0 | 194 |

Continued on next page

Table 6.2: Summary of the simulations used for training data.

| **Simulation Name** | $Ra$ | $Ek$ | $Rm$ | $M_{\text{ratio}}$ | $Q^*$ | **Number of Snapshots** |
|---|---|---|---|---|---|---|
| LEDT016_SM | 30720000 | 0.000063 | 893 | 12 | 0.0 | 194 |
| LEDT019 | 93120000 | 0.000031 | 848 | 15 | 0.0 | 194 |
| LEDT019_SM | 93120000 | 0.000031 | 840 | 16 | 0.0 | 194 |
| LEDT020 | 282240000 | 0.000016 | 805 | 20 | 0.0 | 194 |
| LEDT020_SM | 282240000 | 0.000016 | 772 | 23 | 0.0 | 194 |
| LEDT021 | 19200000 | 0.00010 | 1149 | 7.9 | 0.0 | 194 |
| LEDT021_SM | 19200000 | 0.00010 | 1150 | 8.1 | 0.0 | 194 |
| LEDT022 | 38400000 | 0.00010 | 1015 | 3.5 | 0.0 | 194 |
| LEDT022_SM | 38400000 | 0.00010 | 1002 | 3.8 | 0.0 | 194 |
| LEDT036 | 40320000 | 0.000063 | 1027 | 10 | 0.0 | 194 |
| LEDT036_SM | 40320000 | 0.000063 | 1035 | 9.3 | 0.0 | 194 |
| LEDTF001 | 2000000000 | 0.000070 | 677 | 4.7 | 0.0 | 194 |
| LEDTcase1 | 120000 | 0.0020 | 1116 | 2.6 | 0.0 | 194 |
| LEDTcase2 | 360000 | 0.0010 | 938 | 6.5 | 0.0 | 194 |
| LEDSM001 | 480000 | 0.0010 | 1676 | 3.5 | 0.0 | 194 |
| LEDSM003 | 1750000 | 0.00040 | 2148 | 9.6 | 0.0 | 194 |
| LEDT004_SM | 7680000 | 0.00025 | 1527 | 3.7 | 0.0 | 194 |
| LEDT005 | 12000000 | 0.00020 | 1574 | 3.8 | 0.0 | 194 |
| LEDT005_SM | 12000000 | 0.00020 | 1592 | 3.6 | 0.0 | 194 |
| LEDT006 | 30720000 | 0.00013 | 1694 | 4.2 | 0.0 | 194 |
| LEDT006_SM | 30720000 | 0.00013 | 1690 | 4.0 | 0.0 | 194 |
| LEDT007 | 48000000 | 0.00010 | 1753 | 4.0 | 0.0 | 194 |
| LEDAX053 | 2500000 | 0.00024 | 598 | 4.8 | 0.87 | 142 |
| LEDCJ002 | 200000000 | 0.000020 | 848 | 2.8 | 2.5 | 392 |
| LEDCJ003 | 200000000 | 0.000020 | 828 | 3.6 | 5.3 | 392 |
| LEDCJ006 | 600000000 | 0.000020 | 1496 | 1.3 | 5.3 | 392 |

Continued on next page

Table 6.2: Summary of the simulations used for training data.

| **Simulation Name** | $\boldsymbol{Ra}$ | $\boldsymbol{Ek}$ | $\boldsymbol{Rm}$ | $\boldsymbol{M_{\mathrm{ratio}}}$ | $\boldsymbol{Q^*}$ | **Number of Snapshots** |
|---|---|---|---|---|---|---|
| LEDT012 | 3360000 | 0.00025 | 1134 | 3.5 | 2.5 | 392 |
| LEDT017 | 30720000 | 0.000063 | 910 | 11 | 1.2 | 392 |
| LEDT018 | 30720000 | 0.000063 | 956 | 8.4 | 2.5 | 392 |
| LEDT038 | 120000 | 0.0020 | 801 | 4.6 | 0.13 | 392 |
| LEDT039 | 480000 | 0.0010 | 1018 | 2.9 | 0.13 | 392 |
| LEDT040 | 1920000 | 0.00050 | 1150 | 3.5 | 0.13 | 392 |
| LEDT041 | 7680000 | 0.00025 | 1304 | 3.2 | 0.13 | 392 |
| LEDT041ext | 7680000 | 0.00025 | 1296 | 3.4 | 0.13 | 392 |
| LEDT042 | 12000000 | 0.00020 | 1313 | 4.1 | 0.13 | 392 |
| LEDT043 | 30720000 | 0.00013 | 1401 | 4.1 | 0.13 | 392 |
| LEDT044 | 48000000 | 0.00010 | 1480 | 3.6 | 0.13 | 392 |
| LEDT045 | 3360000 | 0.00025 | 779 | 9.4 | 0.13 | 392 |
| LEDT048 | 450000 | 0.0020 | 744 | 1.0 | 0.13 | 392 |
| LEDT050 | 480000 | 0.0010 | 997 | 3.5 | 0.13 | 392 |
| LEDT052 | 1920000 | 0.00050 | 1163 | 3.2 | 0.13 | 392 |
| LEDT053 | 30720000 | 0.00013 | 1380 | 4.5 | 0.13 | 392 |
| LEDB003 | 500000 | 0.0010 | 408 | 2.1 | 10.7 | 392 |

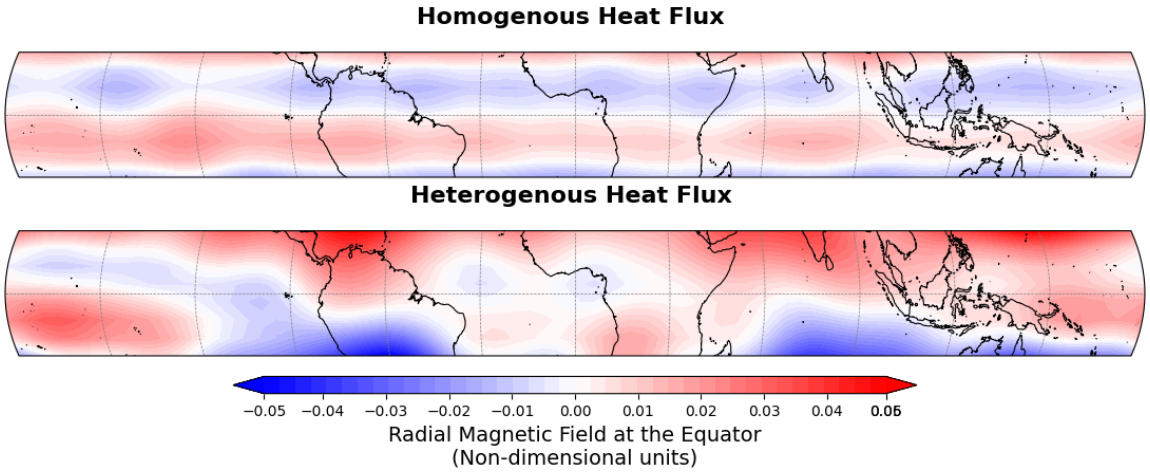



Figure 6.12: Class-averaged radial magnetic field in the equatorial region, truncated at spherical harmonic degree 12. The top panel shows simulations with $Q^* = 0$, and the bottom panel shows simulations with $Q^* \neq 0$.

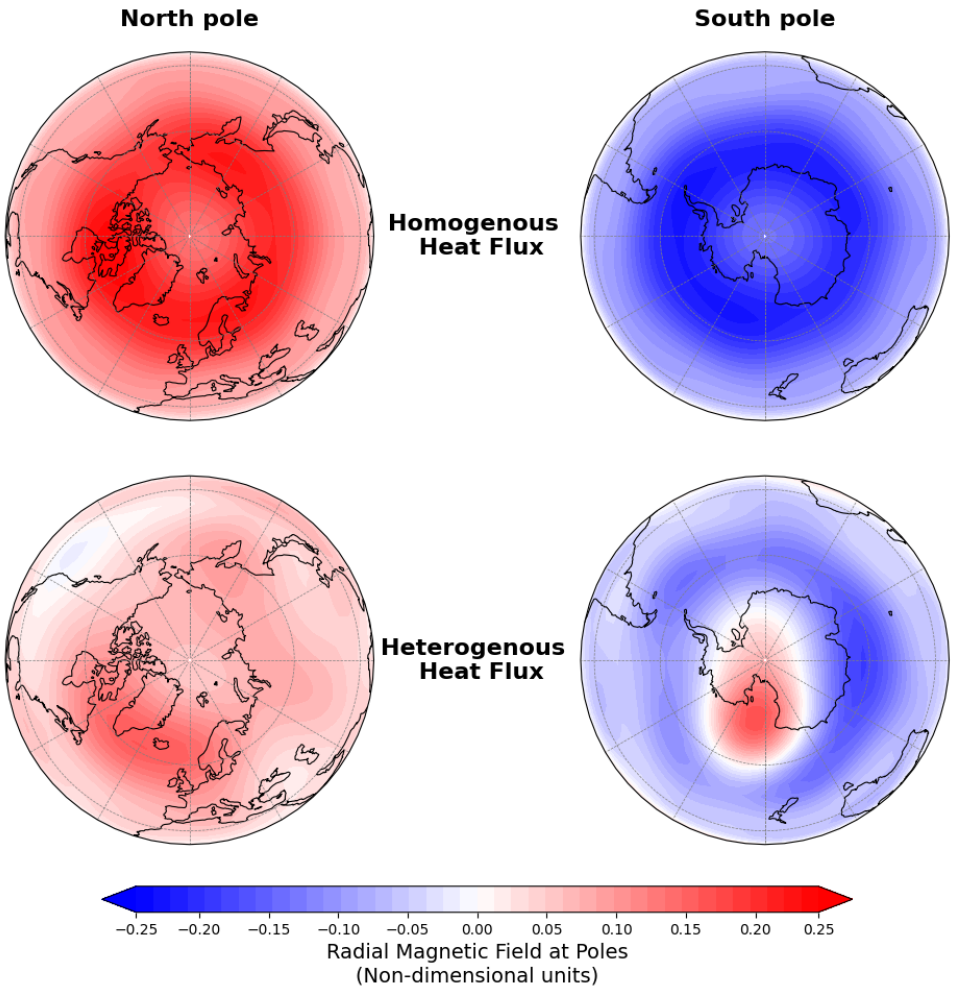



Figure 6.13: Class-averaged radial magnetic field in the polar region, truncated at spherical harmonic degree 12. The top panel shows simulations with $Q^* = 0$, and the bottom panel shows simulations with $Q^* \neq 0$.

The classification errors for the wider simulation set are shown in Table 6.3. When all simulations listed in Table 6.2 were included in the dataset, the training and validation classification errors were 16.9% and 20.1%, respectively. These errors are higher than those obtained in Section 6.3, suggesting that the classifier performs less well when trained on the wider range of simulations. One possible reason for this is that several of the heterogeneous simulations have relatively weak imposed heat-flux heterogeneity, with $Q^* = 0.13$ or $Q^* = 0.87$, compared with the $Q^* > 2$ cases used in Section 6.3. Due to simulation availability, the majority of the simulations in the dataset are low $Q^*$. The effect of such a small $Q^*$ on the resulting radial magnetic field may be weak, making these cases harder to distinguish from homogeneous simulations. To test this, the classifier was re-trained using only heterogenous simulations with $Q^* > 1$ as examples. In this case, the training classification error decreased to 6.6%, while the validation classification error decreased to 12.9%. This improvement suggests that the low-$Q^*$ cases contributed to the poorer performance in the full wider dataset, although the errors remain higher than those obtained using only the LEDCJ simulations in Section 6.3. This has implications for the application to classifiers of the Earth's magnetic field, as the classifier may struggle to distinguish between weak heterogeneity and homogeneity.

| **Training Simulations** | **Classification Error (Training)** | **Classification Error (Validation)** |
|---|---|---|
| All simulations in Table 6.2 | 16.9% | 20.1% |
| All simulations with $Q^* > 1$ | 6.6% | 12.9% |

Table 6.3: Results for wider simulation training data.

The aggregate saliency maps obtained when all simulations were included are shown in the top panel of Figure 6.14. For correctly classified homogeneous cases, the strongest saliency occurs in the equatorial region and at mid-northern latitudes, which is different from the results shown in 6.10. For correctly classified heterogeneous cases, the strongest saliency is found at southern mid-latitudes, close to the boundaries of the Pacific and African LLVPs. Compared with the saliency maps in Section 6.3, these maps are less spatially concentrated, which may reflect the wider range of control parameters and imposed heat-flux patterns included in the training data.

The bottom panel of Figure 6.14 shows the aggregate saliency maps for the case where heterogeneous training examples were restricted to simulations with $Q^* > 1$. The correctly classified homogeneous cases highlight similar regions to those in the top panel of Figure 6.14. In contrast, the correctly classified heterogeneous cases show stronger saliency around the southern Pacific,

similar to the regions highlighted in Figure 6.10, and are less spatially diffuse than those obtained when all heterogeneous simulations were included. This may indicate that, once the weakest heterogeneous cases are removed, magnetic-field structure near the southern Pacific provides a more consistent feature for distinguishing homogeneous and heterogeneous simulations.

The difference between the top and bottom panels of Figures 6.14 may also reflect differences in the imposed heat-flux patterns. Most of the simulations with $Q^* \geq 1$ use a tomographic heat-flux pattern based on Masters et al. (1996), whereas the full heterogeneous dataset includes a broader range of forcing patterns, including those based on Masters et al. (2000). Both datasets also include imposed heat-flux patterns with different spatial scales, with the maximum spherical harmonic degree ranging from 12 to 128.

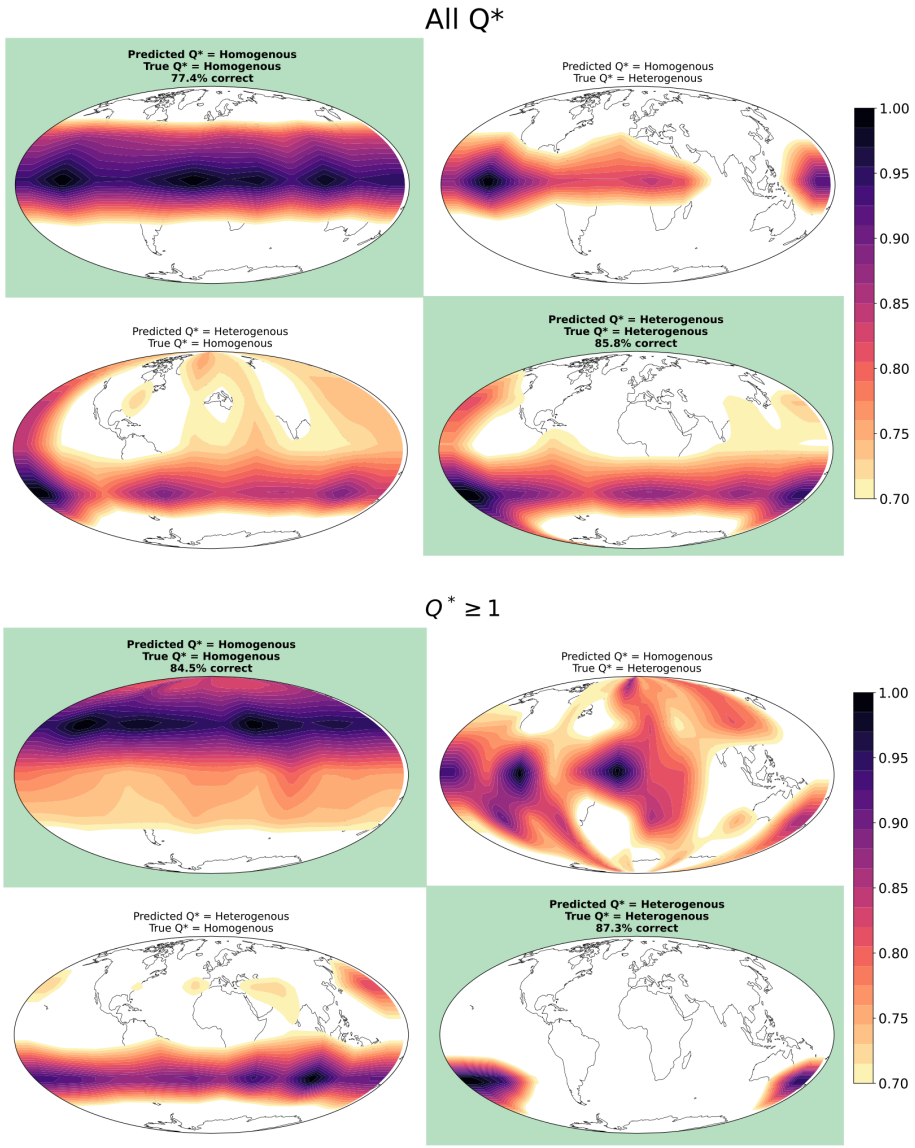



Figure 6.14: [Top] Aggregate saliency maps using the full set of simulations in Table 6.2, shown in a confusion-matrix layout. [Bottom] Aggregate saliency maps using heterogenous simulation where $Q^* > 1$ only. Maps on the diagonal, with a green underlay, correspond to correctly classified samples, while off-diagonal maps correspond to misclassified samples. Darker regions indicate areas that contributed more strongly to the classifier prediction.

## 6.5 Classification using models of the Earth

To test the methodology using observations of the Earth's magnetic field, we used the CHAOS-8 and gufm1 models. CHAOS-8 is a time-dependent model of the near-Earth geomagnetic field spanning 1999–2025, constructed using magnetic measurements from several low-Earth-orbit satellites and ground-based geomagnetic observatories. Its time-dependent internal field is represented up to spherical harmonic degree 20, with the temporal evolution of the spherical harmonic coefficients described using B-splines (Kloss et al., 2026). In contrast, gufm1 is a historical model of the geomagnetic field spanning 1590–1990, constructed from a compilation of historical magnetic observations, including measurements obtained from ship logbooks, together with more recent observatory, survey and satellite data. The model represents the internal magnetic field using a spherical harmonic expansion up to degree 14, with its temporal variation also described using B-splines (Jackson et al., 2000).

To make the Earth models comparable with the simulation data used to train the classifier, the radial magnetic-field snapshots were truncated at spherical harmonic degree $n_{\max} = 12$ and evaluated at the CMB. The fields were then normalised in a similar manner to the simulation data described in Section 6.2, using the average magnetic energy at the CMB calculated across the snapshots from each model, in order to make the Earth–like fields vary on the same scale as the training data. We used epochs 2000.0-2025.0 for the CHAOS-8.5 fields and 1590.0-1990.0 for gufm1. The resulting normalised mean spherical harmonic power spectra for CHAOS-8 and gufm1 are shown in Figure 6.15. These spectra allow the distribution of magnetic-field power across spherical harmonic degree and order to be compared with that of the numerical simulations before the Earth models are passed to the classifier. The power spectra have similar distributions at spherical harmonic degrees below $n = 6$, although CHAOS-8.5 has the greatest power below $n = 4$. Above $n = 6$, the spectra diverge. The simulation spectra, shown in green and purple, remain approximately flat, forming white spectra, whereas the gufm1 and CHAOS-8.5 spectra decrease with increasing spherical harmonic degree, reflecting the regularisation applied in constructing these models.

The methodology was applied to 26 snapshots from CHAOS-8.5 using five classifiers trained identically but initialised with different random seeds. CHAOS-8.5 provides a well-constrained model of the recent geomagnetic field and uses temporal regularisation informed by geodynamo

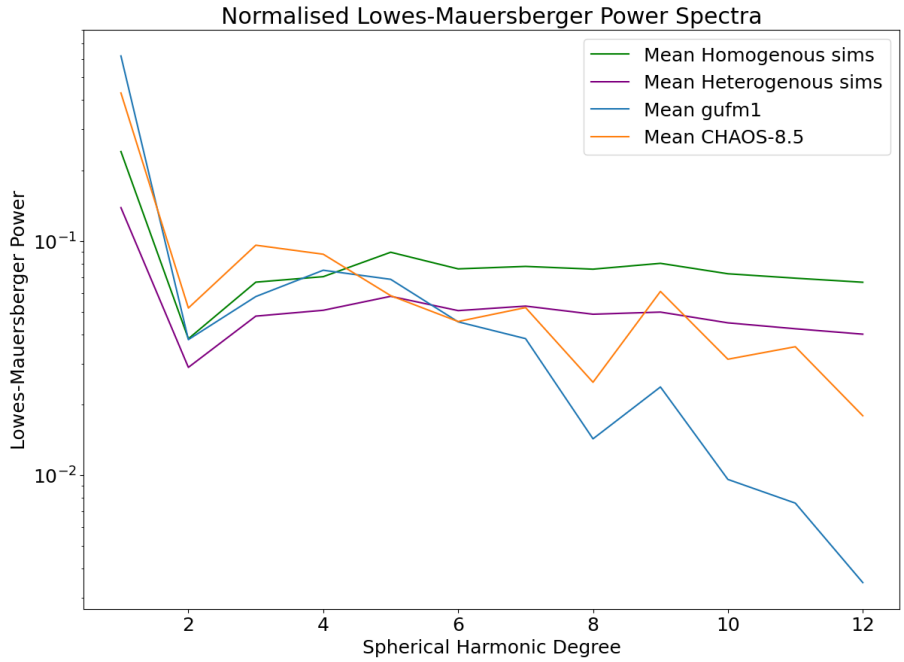



Figure 6.15: Mean normalised Lowes–Mauersberger power spectra for the homogeneous and heterogeneous simulations used for training, and for the CHAOS-8 and gufm1 Earth magnetic-field models.

simulations, making it particularly suitable for comparison with the simulation-based training data. All five classifiers classified every snapshot as corresponding to homogeneous CMB heat flux, showing that the result is consistent across both snapshots and random initialisations.

The aggregate saliency map, shown in Figure 6.16, identifies a region beneath the South Atlantic as contributing most strongly to the classifications. This region lies within the broad latitudinal band highlighted for correctly classified homogeneous simulations in Figure 6.14. The concentration of saliency near the South Atlantic Anomaly, the outline of which is shown in green in Figure 6.16, does not indicate that the anomaly itself is evidence of homogeneous heat flux, but rather that magnetic-field structure in this region contributed strongly to the classification. However, the South Atlantic Anomaly is not well reproduced in geodynamo simulations (Terra-Nova and Amit, 2025). The classifier may therefore be responding to a feature that is not adequately represented in the training data and whose relationship with CMB heat-flux heterogeneity has not been learned. More generally, machine-learning models can perform poorly when applied to data that differ from the distribution represented during training, as reliable generalisation beyond the training distribution remains a major challenge (Goodfellow et al., 2016; Krueger et al., 2021). The classification of the Earth's magnetic field as corresponding to homogeneous CMB heat flux should therefore be interpreted with caution.

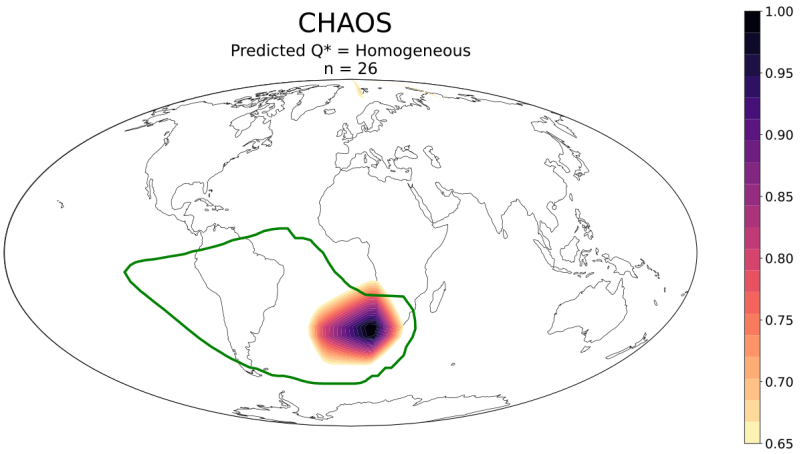



Figure 6.16: Aggregate saliency map for the CHAOS-8.5 classifications, with the boundary of the South Atlantic Anomaly, defined as the contour around the area below 30,000nT at the Earth's surface, shown in green.

The methodology was applied to 41 individual snapshots from gufm1 and one time-averaged field using five classifiers trained identically but initialised with different random seeds. This produced 42 classifications for each classifier. Unlike the CHAOS-8.5 results, the gufm1 classifications were not consistent across either snapshots or random initialisations, as shown in Table 6.4. Although the majority of classifications corresponded to homogeneous CMB heat flux, the results varied considerably between classifiers, with one classifier classifying every input field as homogeneous. This dependence on random initialisation indicates that the gufm1 classification is not robust. This may reflect fragility in the classifier or differences between gufm1 and the magnetic fields represented in the training data. In particular, gufm1 may be less representative of the Earth's field because of the quality and distribution of the historical observations used in its construction, or because it uses stronger regularisation that is not informed by geodynamo simulations. This stronger regularisation is seen in the spectra above $n = 6$ in Figure 6.15.

The aggregate saliency maps for the gufm1 results are shown in Figure 6.17. For classifications corresponding to homogeneous CMB heat flux, the strongest saliency is concentrated in the Northern Hemisphere. In contrast, classifications corresponding to heterogeneous CMB heat flux show stronger contributions from equatorial regions, similar to those highlighted for heterogeneous simulations in the top panel of Figure 6.14. These differences suggest that the classifiers use different spatial features of the gufm1 field when assigning it to the two classes. However, the variation between snapshots and random initialisations means that neither classification can be considered robust.

Table 6.4: Classification results for the individual gufm1 snapshots and the time-averaged field using classifiers initialised with five different random seeds.

| **Year** | **Random seed** | | | | |
|---|---|---|---|---|---|
| | **22** | **56** | **89** | **4875** | **141198** |
| 1590 | Hom | Hom | Hom | Hom | Hom |
| 1600 | Hom | Hom | Hom | Hom | Hom |
| 1610 | Hom | Het | Het | Hom | Hom |
| 1620 | Hom | Het | Het | Hom | Hom |
| 1630 | Hom | Het | Het | Hom | Hom |
| 1640 | Hom | Het | Het | Hom | Hom |
| 1650 | Het | Het | Het | Het | Hom |
| 1660 | Het | Het | Het | Het | Hom |
| 1670 | Het | Het | Het | Het | Hom |
| 1680 | Het | Het | Het | Het | Hom |
| 1690 | Het | Het | Het | Het | Hom |
| 1700 | Het | Het | Het | Het | Hom |
| 1710 | Hom | Hom | Hom | Hom | Hom |
| 1720 | Het | Het | Het | Het | Hom |
| 1730 | Het | Het | Het | Het | Hom |
| 1740 | Het | Het | Het | Het | Hom |
| 1750 | Het | Het | Het | Het | Hom |
| 1760 | Het | Het | Het | Het | Hom |
| 1770 | Het | Het | Het | Het | Hom |
| 1780 | Het | Het | Het | Het | Hom |
| 1790 | Het | Het | Het | Hom | Hom |
| 1800 | Het | Het | Het | Hom | Hom |
| 1810 | Het | Het | Het | Hom | Hom |
| 1820 | Hom | Hom | Hom | Hom | Hom |
| 1830 | Hom | Het | Het | Hom | Hom |
| 1840 | Hom | Het | Het | Hom | Hom |

Continued on next page

Table 6.4: Classification results for the individual gufm1 snapshots and the time-averaged field using classifiers initialised with five different random seeds.

| **Year** | **Random seed** | | | | |
|---|---|---|---|---|---|
| | **22** | **56** | **89** | **4875** | **141198** |
| 1850 | Hom | Het | Het | Hom | Hom |
| 1860 | Hom | Het | Het | Het | Hom |
| 1870 | Hom | Het | Het | Hom | Hom |
| 1880 | Hom | Het | Het | Hom | Hom |
| 1890 | Hom | Hom | Het | Hom | Hom |
| 1900 | Hom | Hom | Het | Hom | Hom |
| 1910 | Hom | Hom | Het | Hom | Hom |
| 1920 | Hom | Hom | Hom | Hom | Hom |
| 1930 | Hom | Hom | Hom | Hom | Hom |
| 1940 | Hom | Hom | Hom | Hom | Hom |
| 1950 | Hom | Hom | Hom | Hom | Hom |
| 1960 | Hom | Hom | Hom | Hom | Hom |
| 1970 | Hom | Hom | Hom | Hom | Hom |
| 1980 | Hom | Hom | Hom | Hom | Hom |
| 1990 | Hom | Hom | Hom | Hom | Hom |
| TAF | Hom | Hom | Het | Hom | Hom |

Table 6.4: Classification results for the individual gufm1 snapshots and the Time-Averaged Field (TAF) using classifiers initialised with five different random seeds.

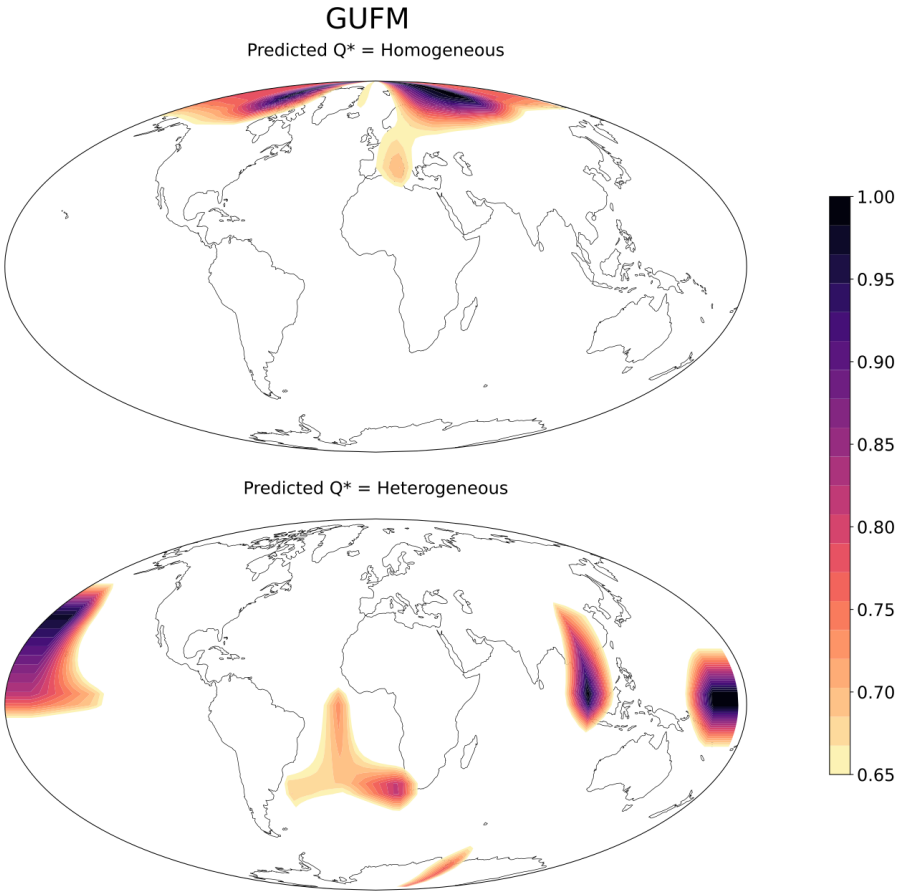



Figure 6.17: [Top] Aggregate saliency map for snapshots classified as homogenous, [Bottom] Aggregate saliency map for snapshots classified as heterogenous.

## 6.6 Discussion and Conclusions

The controlled LEDCJ simulations show that the morphology of the radial magnetic field contains information about imposed CMB heat-flux heterogeneity. The classifier distinguished between homogeneous, weakly heterogeneous and strongly heterogeneous cases with training and validation errors of 1.5% and 4.9%, respectively. The lower error obtained for the two-class problem indicates that identifying the presence of heterogeneity is easier than distinguishing between different amplitudes. Nevertheless, the three-class results show that, when the imposed heat-flux pattern is fixed, the radial magnetic field contains sufficient information to separate different amplitudes of heterogeneity.

The sensitivity tests provide further insight into the magnetic-field information used by the classifier. The substantial increase in classification error when the field was truncated below spherical harmonic degree $n = 4$ shows that non-dipolar structure is important for distinguishing between the heat-flux classes. The improved performance obtained using average scaling rather than scaling each image independently also indicates that relative field amplitude contributes to the classification. The regional tests show that most of the useful information is concentrated at equatorial and mid-latitudes, while the polar regions alone produce substantially higher classification errors. This is consistent with the saliency maps, which repeatedly highlight regions near the Pacific and African LLVPs, particularly beneath the southern Pacific. Although the polar regions are highlighted in some saliency maps, their poor performance when used in isolation shows that they are not sufficient for accurate classification. The classifier therefore appears to use a combination of relative field strength, non-dipolar structure and regional magnetic-field morphology, rather than relying on a single diagnostic feature.

The wider simulation set provides a more demanding test because the simulations vary not only in $Q^*$, but also in their control parameters, dynamo regime and imposed heat-flux pattern. The validation error increased to 20.1% for the full dataset, corresponding to an accuracy of 79.9%, but decreased to 12.9% when heterogeneous cases with $Q^* < 1$ were removed. This suggests that stronger imposed heterogeneity produces a clearer magnetic signature, while weakly heterogeneous cases may be difficult to distinguish from the natural variability of homogeneous simulations. This is consistent with the saliency maps, which are more diffuse for the full dataset but become concentrated around the southern Pacific when only cases with $Q^* > 1$ are

included. Magnetic-field structure near the Pacific LLVP therefore appears to provide a more consistent distinction between homogeneous and heterogeneous simulations when the imposed heterogeneity is sufficiently strong.

However, the effect of heat-flux heterogeneity cannot be fully separated from other differences between the simulations. Changes in $Q^*$ are accompanied by variations in the control parameters and imposed heat-flux pattern, meaning that the classifier may respond to features associated with these differences rather than to $Q^*$ alone. The limited number of strongly heterogeneous simulations also meant that the wider dataset could only be treated as a two-class problem, preventing different amplitudes of heterogeneity from being compared directly. Nevertheless, the classifier performed above chance across this more varied dataset, showing that the radial magnetic field retains information about sufficiently strong heat-flux heterogeneity when other simulation properties vary.

Application to models of the Earth's magnetic field tested whether the classifier could generalise beyond the simulations used for training. All CHAOS-8.5 snapshots were classified as corresponding to homogeneous CMB heat flux by all five classifier initialisations. This consistency suggests that CHAOS-8.5 contains features that the classifier associates strongly with the homogeneous simulation class. This is broadly consistent with the power spectra in Figure 6.15, since CHAOS-8.5 and the homogeneous simulations have similar power distributions at low spherical harmonic degrees. However, the saliency map shows that magnetic-field structure near the South Atlantic Anomaly contributed strongly to the prediction. Since this feature is not well reproduced in the training simulations, it is unclear whether the classifier has identified a heat-flux-related signature or has responded to magnetic structure that was not adequately represented during training. The CHAOS-8.5 classification cannot therefore be interpreted as direct evidence that the Earth's CMB heat flux is homogeneous.

The gufm1 classifications were less consistent. Although most predictions corresponded to homogeneous CMB heat flux, the predicted class varied between epochs and random initialisations, and one classifier classified every input as homogeneous. This classifier achieved a validation error similar to those of the other classifiers, indicating that its behaviour cannot be attributed to poor training. The decreasing power in gufm1 above spherical harmonic degree $n = 6$, shown in Figure 6.15, indicates that its smaller-scale field structure differs from that of the simulations. This may partly reflect differences in the observations and regularisation used to

construct gufm1 and CHAOS-8.5. The gufm1 saliency maps also show that the two predicted classes depend on different spatial features. Homogeneous classifications primarily highlight the Northern Hemisphere, whereas heterogeneous classifications place greater emphasis on equatorial regions similar to those identified in the heterogeneous simulations in Section 6.4. The dependence on epoch, random initialisation and geomagnetic field model means that the gufm1 classifications do not provide a robust inference of the Earth's CMB heat-flux structure.

The contrasting CHAOS-8.5 and gufm1 results show that the predicted heat-flux class is sensitive to differences in field morphology, spectral content and regularisation. A more representative training ensemble is therefore required before the methodology can be applied reliably to the Earth. This should include a wider range of dynamo regimes, heat-flux heterogeneity amplitudes and spatial forcing patterns, with individual parameters varied systematically where possible. Such an ensemble would help separate magnetic-field changes caused by heat-flux heterogeneity from those associated with the broader dynamo regime. A probabilistic approach could also provide a more informative measure of confidence than a single classification, particularly where predictions vary between random initialisations or geomagnetic field models.

Overall, the results show that simulated radial magnetic fields contain classifiable information about imposed CMB heat-flux heterogeneity, particularly when the forcing pattern is fixed and its amplitude is sufficiently strong. However, the disagreement between the CHAOS-8.5 and gufm1 results means that this conclusion cannot yet be extended reliably to the Earth. The Earth-model results should therefore be treated as demonstrating the current limitations of transferring a classifier trained on numerical simulations to geomagnetic field models, rather than as evidence for either homogeneous or heterogeneous CMB heat flux.

## 6.7 Summary

This study presents a convolutional autoencoder classifier for testing whether CMB heat-flux heterogeneity can be inferred from snapshots of the radial magnetic field. We first apply the method using the controlled LEDCJ simulations, where the classifier distinguishes between homogeneous and heterogeneous heat-flux cases with low error, less than 10% with the shuffled partition approach, and can also separate different amplitudes of heterogeneity when the imposed heat-flux pattern is fixed. The sensitivity tests show that classification depends on both relative field strength and non-dipolar magnetic structure, with performance decreasing when the field is truncated below spherical harmonic degree $n = 4$. Regional tests and saliency maps indicate that the most useful information is concentrated at equatorial and mid-latitudes, particularly near the Pacific and African LLVPs, while the polar regions alone are insufficient. Applying the method to a wider set of simulations increases the classification error, showing that the magnetic imprint of heat-flux heterogeneity is harder to isolate when the dynamo regime, boundary conditions, and magnetic-field morphology vary between simulations. Removing the weakly heterogeneous cases improves the classification accuracy, suggesting that, for small values of $Q^* < 1$, the magnetic imprint may be too weak, too variable, or too similar to the variability found in homogeneous simulations. Application to models of the Earth's magnetic field gives different outcomes: CHAOS-8.5 is consistently classified as corresponding to homogeneous CMB heat flux, whereas the gufm1 classifications vary between epochs and random initialisations. Overall, these results show that simulated CMB magnetic fields contain classifiable information about imposed heat-flux heterogeneity. However, the differences between the CHAOS-8.5 and gufm1 classifications show that further work is needed before results obtained by applying a simulation-trained classifier to the Earth's magnetic field can be considered reliable.

# Appendix

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| LEDCJ001 | Mound, J. E., Davies, C. J. (2023) Nature Geoscience | 0.0 | Homog |
| LEDSM002 | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT001 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT002 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT002_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT003 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| LEDT003_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT004 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT008 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT008_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT009 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT009_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT010 | Nakagawa, T., | | |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| | Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT010_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT013 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT013_SM | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT014_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT015 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT015_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| LEDT016 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT016_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT019 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT019_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT020 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT020_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT021 | Nakagawa, T., | | |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| | Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT021_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT022 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT022_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT036 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT036_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDTF001 | Frasson, T., Davies, C. J., | | |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| | Biggin, A. J. (2026, submitted) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDTcase1 | Christensen, U. R., Aubert, J., Cardin, P., Dormy, E., Gibbons, S., Glatzmaier, G. A., Grote, E., Honkura, Y., Jones, C., Kono, M., Matsushima, M., Sakuraba, A., Takahashi, F., Tilger, A., Wicht, J., Zhang, K. (2001) Physics of the Earth and Planetary Interiors | 0.0 | Homog |
| LEDTcase2 | Christensen, U. R., Aubert, J., Cardin, P., Dormy, E., Gibbons, S., Glatzmaier, G. A., Grote, E., Honkura, Y., Jones, C., Kono, M., Matsushima, M., Sakuraba, A., Takahashi, F., Tilger, A., Wicht, J., | | |
| | Zhang, K. (2001) Physics of the Earth and Planetary Interiors | 0.0 | Homog |
| LEDSM001 | Mason, S. J., Davies, C. J., Clarke, A. T., | | |
| | Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| LEDSM003 | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT004_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT005 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT005_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT006 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDT006_SM | Mason, S. J., Davies, C. J., Clarke, A. T., Constable, C. G. (2024) Earth and Planetary Science Letters | 0.0 | Homog |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| LEDT007 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.0 | Homog |
| LEDAX053 | Unpublished, formulation found in Davies, C. J., Constable, C. G. (2014); Earth and Planetary Science Letters. | 0.87 | Masters (1996), $n_{\max}$ =12 |
| LEDCJ002 | Mound, J. E., Davies, C. J. (2023) Nature Geoscience | 2.46 | Masters (1996), $n_{\max}$ =12 |
| LEDCJ003 | Mound, J. E., Davies, C. J. (2023) Nature Geoscience | 5.35 | Masters (1996), $n_{\max}$ =12 |
| LEDCJ006 | Mound, J. E., Davies, C. J. (2023) Nature Geoscience | 5.35 | Masters (1996), $n_{\max}$ =12 |
| LEDT012 | Nakagawa, T., | | |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| | Davies, C. J. (2022) Earth and Planetary Science Letters | 2.50 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT017 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 1.25 | Masters (2000), $n_{\max}$ =128 (Codensity) |
| LEDT018 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 2.50 | Masters (2000), $n_{\max}$ =128 (Codensity) |
| LEDT038 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT039 | Nakagawa, T., | | |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| | Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT040 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT041 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =128 (Codensity) |
| LEDT041ext | Unpublished, formulation described in Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =128 (Codensity) |
| LEDT042 | Nakagawa, T., | | |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| | Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =128 (Codensity) |
| LEDT043 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =128 (Codensity) |
| LEDT044 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =128 (Codensity) |
| LEDT045 | Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT048 | Nakagawa, T., | | |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| | Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT050 | Unpublished, formulation described in Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT052 | Unpublished, formulation described in Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |
| LEDT053 | Unpublished, formulation described in Nakagawa, T., Davies, C. J. (2022) Earth and Planetary Science Letters | 0.13 | Masters (2000), $n_{\max}$ =96 (Codensity) |

Continued on next page

Table 6.5: Reference and forcing metadata for the simulations listed in Table 6.2.

| Name | Reference | $Q^*$ | $Q^*$ pattern |
|---|---|---|---|
| LEDB003 | Unpublished, formulation described in Willis, A. P., Sreenivasan, B., Gubbins, D. (2007); Physics of the Earth and Planetary Interiors | 10.72 | Masters (1996), $n_{\max}=12$ |

# References


Amit, H., Aubert, J., Hulot, G., 2010. Stationary, oscillating or drifting mantle-driven geomagnetic flux patches? Journal of Geophysical Research: Solid Earth 115. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2009JB006542, doi:https://doi.org/10.1029/2009JB006542, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2009JB006542.

Amit, H., Korte, M., Aubert, J., Constable, C., Hulot, G., 2011. The time-dependence of intense archeomagnetic flux patches. Journal of Geophysical Research: Solid Earth 116. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011JB008538, doi:https://doi.org/10.1029/2011JB008538, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2011JB008538.

Aubert, J., Gastine, T., Fournier, A., 2017. Spherical convective dynamos in the rapidly rotating asymptotic regime. Journal of Fluid Mechanics 813, 558–593. doi:10.1017/jfm.2016.789.

Bank, D., Koenigstein, N., Giryes, R., 2023. Autoencoders. Springer International Publishing, Cham. pp. 353–374. URL: https://doi.org/10.1007/978-3-031-24628-9_16, doi:10.1007/978-3-031-24628-9_16.

Biggin, A.J., Davies, C.J., Mound, J.E., Lloyd, S.J., Engbers, Y.E., Thallner, D., Clarke, A.T., Bono, R.K., 2026. Mantle heterogeneity influenced earth's ancient magnetic field. Nature Geoscience doi:10.1038/s41561-025-01910-1.

Bloxham, J., Gubbins, D., 1987. Thermal core– mantle interactions. Nature 325, 511–513. URL: https://doi.org/10.1038/325511a0, doi:10.1038/325511a0.

Dannberg, J., Gassmöller, R., Thallner, D., LaCombe, F., Sprain, C., 2024. Changes in core–mantle boundary heat flux patterns throughout the supercontinent cycle. Geophysical Journal International 237, 1251–1274. doi:10.1093/gji/ggae075.

Dobson, J.E., 2024. On the confusion matrix. Configurations 32, 331–350. doi:10.1353/con.2024.a942087.

Doshi-Velez, F., Kim, B., 2017. Towards a rigorous science of interpretable machine learning. URL: https://arxiv.org/abs/1702.08608, arXiv:1702.08608.

Frasson, T., Labrosse, S., Nataf, H.C., Coltice, N., Flament, N., 2024. On the impact of true polar wander on heat flux patterns at the core–mantle boundary. Solid Earth 15, 617–637. URL: https://se.copernicus.org/articles/15/617/2024/, doi:10.5194/se-15-617-2024.

Frasson, T., Schaeffer, N., Nataf, H.C., Labrosse, S., 2025. Geomagnetic dipole stability and zonal flows controlled by mantle heat flux heterogeneities. Geophysical Journal International 240, 1481–1504. URL: https://doi.org/10.1093/gji/ggae457, doi:10.1093/gji/ggae457, arXiv:https://academic.oup.com/gji/article-pdf/240/3/1481/61290993/ggae457.pdf.

Frost, D.A., Avery, M.S., Buffett, B.A., Chidester, B.A., Deng, J., Dorfman, S.M., Li, Z., Liu, L., Lv, M., Martin, J.F., 2022. Multidisciplinary constraints on the thermal-chemical boundary between earth's core and mantle. Geochemistry, Geophysics, Geosystems 23, e2021GC009764. doi:10.1029/2021GC009764.

Ganotra, R., Gupta, S., Dora, S., 2024. A joint autoencoder and classifier deep neural network for ad and mci classification. International Journal of Imaging Systems and Technology 34, e23054. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/ima.23054, doi:https://doi.org/10.1002/ima.23054, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/ima.23054.

Glatzmaier, G.A., Coe, R.S., Hongre, L., Roberts, P.H., 1999. The role of the Earth's mantle in controlling the frequency of geomagnetic reversals. Nature 401, 885–890. doi:10.1038/44776.

Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. MIT Press. http://www.deeplearningbook.org.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, IEEE. pp. 770–778.

Huber, P.J., 1964. Robust estimation of a location parameter. The Annals of Mathematical Statistics 35, 73–101. doi:10.1214/aoms/1177703732.

Huguet, L., Amit, H., Alboussière, T., 2018. Geomagnetic dipole changes and upwelling/downwelling at the top of the earth's core. Frontiers in Earth Science Volume 6 - 2018. URL: https://www.frontiersin.org/journals/earth-science/articles/10.3389/feart.2018.00170, doi:10.3389/feart.2018.00170.

Jackson, A., Jonkers, A.R.T., Walker, M.R., 2000. Four centuries of geomagnetic secular variation from historical records. Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 358, 957–990. doi:10.1098/rsta.2000.0569.

Jordan, Mitchell, 2015. Machine learning: Trends, perspectives, and prospects. Science 349, 255–260. doi:10.1126/science.aaa8415.

Kloss, C., Finlay, C.C., Olsen, N., Tøffner-Clausen, L., Gillet, N., Grayver, A., 2026. The CHAOS-8 Geomagnetic Field Model. Earth, Planets and Space 78, 21. URL: https://doi.org/10.1186/s40623-025-02352-0, doi:10.1186/s40623-025-02352-0.

Koelemeijer, P., 2021. Toward Consistent Seismological Models of the Core–Mantle Boundary Landscape. American Geophysical Union (AGU). chapter 9. pp. 229–255. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/9781119528609.ch9, doi:https://doi.org/10.1002/9781119528609.ch9, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/9781119528609.ch9.

Kozmin, A., Kalashev, O., Chernenko, A., Redyuk, A., 2025. Semi-supervised learned autoencoder for classification of events in distributed fibre acoustic sensors. Sensors 25, 3730. doi:10.3390/s25123730.

Krueger, D., Caballero, E., Jacobsen, J.H., Zhang, A., Binas, J., Zhang, D., Le Priol, R., Courville, A., 2021. Out-of-distribution generalization via risk extrapolation, in: Proceedings of the 38th International Conference on Machine Learning.

Landeau, M., Fournier, A., Nataf, H.C., Cébron, D., Schaeffer, N., 2022. Sustaining earth's magnetic dynamo. Nature Reviews Earth & Environment 3, 255–269. doi:10.1038/s43017-022-00264-1.

Linardatos, P., Papastefanopoulos, V., Kotsiantis, S., 2021. Explainable ai: A review of machine learning interpretability methods. Entropy 23. URL: https://www.mdpi.com/1099-4300/23/1/18, doi:10.3390/e23010018.

Loshchilov, I., Hutter, F., 2017. Fixing weight decay regularization in adam. CoRR abs/1711.05101. URL: http://arxiv.org/abs/1711.05101, arXiv:1711.05101.

Lowes, F.J., 1974. Spatial power spectrum of the main geomagnetic field, and extrapolation to the core. Geophysical Journal International 36, 717–730. URL: https://doi.org/10.1111/j.1365-246X.1974.tb00622.x, doi:10.1111/j.1365-246X.1974.tb00622.x, arXiv:https://academic.oup.com/gji/article-pdf/36/3/717/1576102/36-3-717.pdf.

Mao, J., Guo, S., Yin, X., Chang, Y., Nie, B., Wang, Y., 2025. Medical supervised masked autoencoder: Crafting a better masking strategy and efficient fine-tuning schedule for medical image classification. Applied Soft Computing 169, 112536. URL: https://www.sciencedirect.com/science/article/pii/S1568494624013103, doi:https://doi.org/10.1016/j.asoc.2024.112536.

Masters, G., Laske, G., Bolton, H., Dziewonski, A., 2000. The Relative Behavior of Shear Velocity, Bulk Sound Speed, and Compressional Velocity in the Mantle: Implications for Chemical and Thermal Structure. American Geophysical Union (AGU). pp. 63–87. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/GM117p0063, doi:https://doi.org/10.1029/GM117p0063, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/GM117p0063.

Masters, T.G., Johnson, S., Laske, G., Bolton, H., 1996. A shear - velocity model of the mantle. Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 354, 1385–1411. doi:10.1098/rsta.1996.0054.

Mosca, I., Cobden, L., Deuss, A., Ritsema, J., Trampert, J., 2012. Seismic and mineralogical structures of the lower mantle from probabilistic tomography. Journal of Geophysical Research: Solid Earth 117. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2011JB008851, doi:https://doi.org/10.1029/2011JB008851, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2011JB008851.

Mound, J., Davies, C., Rost, S., Aurnou, J., 2019. Regional stratification at the top of Earth's core due to core–mantle boundary heat flux variations. Nature Geoscience 12, 575–580. URL: https://doi.org/10.1038/s41561-019-0381-z, doi:10.1038/s41561-019-0381-z.

Mound, J.E., Davies, 2023. Longitudinal structure of earth's magnetic field controlled by lower mantle heat flow. Nature Geoscience 16, 380–385. doi:10.1038/s41561-023-01148-9.

Nakagawa, T., Davies, C.J., 2022. Combined dynamical and morphological characterisation of geodynamo simulations. Earth and Planetary Science Letters 594, 117752. URL: https://www.sciencedirect.com/science/article/pii/S0012821X22003880, doi:https://doi.org/10.1016/j.epsl.2022.117752.

Naskar, S., Mound, J.E., Davies, C.J., Clarke, A.T., 2026. Thermochemical models of outer core convection with heterogeneous core-mantle boundary heat flux. Journal of Studies of Earth's Deep Interior Volume 2. URL: https://jsedi.episciences.org/17084, doi:10.46298/jsedi.17084.

Panovska, S., Korte, M., Constable, C.G., 2019. One hundred thousand years of geomagnetic field evolution. Reviews of Geophysics 57, 1289–1337. doi:10.1029/2019RG000656.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. Pytorch: An imperative style, high-performance deep learning library. URL: https://arxiv.org/abs/1912.01703, arXiv:1912.01703.

Rogers, H., Davies, C., Constable, C., Mason, S., Frasson, T., 2026. Using paleomagnetic observations and geodynamo simulations to assess earth's magnetic field morphology and variability. Journal of Studies of the Earth's Deep Interior 2. URL: https://doi.org/10.5281/zenodo.21821109, doi:10.5281/zenodo.21821109.

Saeedi, J., Giusti, A., 2023. Semi-supervised visual anomaly detection based on convolutional autoencoder and transfer learning. Machine Learning with Applications 11, 100451. URL: https://www.sciencedirect.com/science/article/pii/S266682702300004X, doi:https://doi.org/10.1016/j.mlwa.2023.100451.

Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization, in: Proceedings of the IEEE International Conference on Computer Vision (ICCV).

Sola, J., Sevilla, J., 1997. Importance of input data normalization for the application of neural networks to complex industrial problems. IEEE Transactions on Nuclear Science 44, 1464–1468. doi:10.1109/23.589532.

Su, S.Z., Liu, Z.H., Xu, S.P., Li, S.Z., Ji, R., 2015. Sparse auto-encoder based feature learning for human body detection in depth image. Signal Processing 112, 43–52. URL: https://www.sciencedirect.com/science/article/pii/S0165168414005052, doi:https://doi.org/10.1016/j.sigpro.2014.11.003. signal Processing and Learning Methods for 3D Semantic Analysis.

Terra-Nova, F., Amit, H., 2025. Necessary core-mantle boundary heat flux patterns for recovering the latitude of the south atlantic anomaly. Journal of Geophysical Research: Solid Earth 130, e2025JB031193. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2025JB031193, doi:https://doi.org/10.1029/2025JB031193, arXiv:https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2025JB031193. e2025JB031193 2025JB031193.

van Tent, R., Deuss, A., Kaneshima, S., Thomas, C., 2020. The signal of outermost-core stratification in body-wave and normal-mode data. Geophysical Journal International 223, 1338–1354. URL: https://doi.org/10.1093/gji/ggaa368, doi:10.1093/gji/ggaa368, arXiv:https://academic.oup.com/gji/article-pdf/223/2/1338/33763410/ggaa368.pdf.

Whaler, K.A., 1980. Does the whole of the earth's core convect? Nature 287, 528–530. doi:10.1038/287528a0.

Xing, C., Ma, L., Yang, X., 2016. Stacked denoise autoencoder based feature extraction and classification for hyperspectral images. Journal of Sensors 2016, 3632943. URL: https://onlinelibrary.wiley.com/doi/abs/10.1155/2016/3632943, doi:https://doi.org/10.1155/2016/3632943, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1155/2016/3632943.

Zhao, C., Garnero, E.J., McNamara, A.K., Schmerr, N., Carlson, R.W., 2015. Seismic evidence for a chemically distinct thermochemical reservoir in earth's deep mantle beneath hawaii. Earth and Planetary Science Letters 426, 143–153. URL: https://www.sciencedirect.com/science/article/pii/S0012821X15003684, doi:https://doi.org/10.1016/j.epsl.2015.06.012.

Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2921–2929. doi:10.1109/CVPR.2016.319.

# Chapter 7

# Conclusions

This thesis develops and applies machine-learning methods to investigate the Earth's magnetic field. It uses these methods to interpret geophysical processes at the top of the core and the core–mantle boundary.

First, in Chapter 4, we developed a novel regional core-flow inversion methodology using Physics-Informed Neural Networks. The method was designed to infer core-surface flows using only the magnetic field and secular variation within individual regions, while incorporating the tangentially geostrophic flow constraint during training. After validation using synthetic flows, the methodology was applied to CHAOS-8.1 to investigate whether features previously identified using global inversions could also be recovered from regional information. The novel contribution of this chapter was the development and validation of a flexible PINN framework for regional core-flow inversion.

In Chapter 5, we used the core-flow inversion methods developed in the previous chapter to forecast geomagnetic secular variation. Regional and global flow models were constructed from CHAOS-7.18 and used in a steady-flow advection scheme to predict changes in the magnetic field. After validation through hindcasting over the IGRF-13 period, the regional methodology was used to produce a candidate secular-variation model for IGRF-14, covering 2025–2030. The novel contribution of this chapter was the application of the regional PINN-based core-flow inversion framework to secular-variation forecasting, which showed improved hindcast performance compared with the global approach and the IGRF-13 candidate models.

Finally, in Chapter 6, we developed a convolutional autoencoder classifier to test whether CMB

heat-flux heterogeneity could be inferred from snapshots of the radial magnetic field. The method was first validated using a controlled set of geodynamo simulations before being applied to a wider range of simulations and to the CHAOS-8.5 and gufm1 models of the Earth's magnetic field. Sensitivity tests and saliency maps were used to determine which spatial and spectral features of the magnetic field contributed most strongly to the classifications. The novel contribution of this chapter was the development and validation of an explainable machine-learning framework for identifying signatures of CMB heat-flux heterogeneity in radial magnetic-field images.

This discussion revisits the research questions introduced in Chapter 3 and considers how the findings of each research chapter contribute to answering them. For Chapters 4 and 5, the discussion also considers developments made following the publication of Shakespeare-Rees et al. (2025) and Shakespeare-Rees et al. (2026). Finally, the main outcomes and wider implications of the thesis are summarised, together with priorities for future research.

## 7.1 Can regional core-flow inversions recover the large-scale structures found in global flow models?

In Chapter 4, we developed regional core-flow models that were trained independently using only the magnetic field and secular variation within their respective regions. These core flow models were then combined in a mosaic to compare to global flow models. Despite having no knowledge of the secular variation outside these regions, the combined results reproduced previously identified large-scale structures, including westward flow beneath the Atlantic, variable flow beneath Indonesia and Central America, and the eccentric planetary gyre (Pais and Jault, 2008; Gillet et al., 2022; Kloss and Finlay, 2019). The recovery of mostly continuous flow structures across the regional boundaries indicates that the gyre is supported by the magnetic-field data, rather than being imposed by the global representation. The connection between the high- and mid-latitude flows also suggests that at least part of the high-latitude secular variation is associated with the gyre. This does not exclude the contribution of high-latitude waves proposed by Buffett et al. (2016), but provides independent evidence for a flow component.

The regional methodology also recovered time-dependent changes in smaller-scale flow features. These included repeated changes in equatorial azimuthal flow direction, a recent change beneath the eastern Pacific, and a decrease in the intensity of the northern high-latitude jet after

2017. The physical origins of these changes remain uncertain. The equatorial variations may reflect wave-driven behaviour, whereas changes in the high-latitude jet may arise from evolving dynamics near the inner-core tangent cylinder (Gillet et al., 2022; Finlay et al., 2023). Grüne et al. (2025) similarly identified periodic changes in azimuthal flow acceleration at low latitudes, which they interpreted as signatures of propagating and standing waves. The polarity changes in these acceleration features occurred at intervals similar to those between recent geomagnetic jerks, suggesting a possible connection between wave-driven changes in core-surface flow and geomagnetic jerks.

The large-scale flow features recovered by the regional models in Chapter 4 were also identified in Chapter 5, despite the latter not imposing the tangentially geostrophic flow assumption. This suggests that these features may be robust to changes in the flow assumption. However, only the tangentially geostrophic assumption was implemented in Chapter 4, and additional physical constraints must be incorporated into the PINN framework before this can be tested more thoroughly. Further work could incorporate alternative flow assumptions to determine how the choice of constraint affects regional flow estimates. The method of implementation would depend on the assumption considered. A toroidal-only constraint could be imposed straightforwardly by adding a loss term that penalises departures of the poloidal flow, defined in Equation (4.8), from zero. Implementing quasi-geostrophy would be more complex because of the required equatorial symmetry. One possible approach would be to train paired regional inversions in the Northern and Southern Hemispheres, with an additional loss term penalising departures from the appropriate symmetry between them. However, this might increase the complexity of the loss landscape, and the appropriate weighting of the symmetry constraint would require careful investigation.

In Chapter 4, we also chose to implement the flow assumptions as ‘soft constraints’, rather than as ‘hard constraints’ that are satisfied exactly through the construction of the network output, as discussed in Section 2.2.2. Hard constraints have previously been implemented in PINNs by constructing trial functions that satisfy prescribed boundary conditions, parameterising the network outputs to obey conservation laws, or projecting predictions onto a constrained solution space (Lagaris et al., 1998; Sukumar and Srivastava, 2022; Berrone et al., 2023; Horne et al., 2025). These approaches guarantee satisfaction of the imposed condition and avoid the need to balance its residual against other loss terms, but the required parameterisation can be difficult

to construct and may restrict the range of solutions available to the network (Luong et al., 2024; Sheikholeslami et al., 2025). In the present problem, tangential geostrophy could be imposed exactly through the form of the induction equation given in Equation (1.34). However, this would prevent the network from recovering flows that depart from tangential geostrophy, even where the assumption may not be valid. Implementing it as a soft constraint instead allows the model to balance agreement with the secular variation against agreement with the assumed flow physics, while the remaining physical residual can indicate where the assumption may break down. Although a soft constraint does not remove the inversion's non-uniqueness, enforcing tangential geostrophy exactly would only restrict the possible solutions, rather than guarantee uniqueness. The soft formulation was therefore preferred because it does not assume that tangential geostrophy holds exactly at all locations and times.

The secular-variation and flow-constraint losses in Chapter 4 were balanced using a fixed weighting factor determined from their typical magnitudes. However, the relative magnitudes and gradients of these terms can change during training, meaning that a weighting chosen initially may not remain appropriate throughout the optimisation. Adaptive approaches have therefore been developed in which the loss weights are updated during training. For example, Wang et al. (2021) used gradient statistics to balance the contributions of different loss terms, while McClenny and Braga-Neto (2023) introduced trainable pointwise weights that direct the network towards regions with larger residuals. Bischof and Kraus (2025) compared several adaptive strategies and proposed relative loss balancing with random lookback, which adjusts the weights using changes in the individual losses. These methods can reduce manual tuning and prevent one loss term from dominating the optimisation, but they also introduce additional calculations and choices governing how frequently and rapidly the weights are updated. The fixed weighting used here was sufficient to bring the two losses to similar magnitudes by the end of training, but adaptive weighting could provide a more flexible approach, particularly where the balance between fitting the secular variation and satisfying the flow constraint changes during optimisation.

Overall, the results show that regional core-flow inversions can recover the principal large-scale structures identified in global flow models. However, CHAOS-8.1 is itself a global geomagnetic field model, meaning that the magnetic field and secular variation supplied to each regional inversion remain influenced by observations from elsewhere. Future work could therefore apply

the methodology to independently constructed regional geomagnetic field estimates, such as those obtained using SOLA (Subtractive Optimally Localised Averages, Hammer and Finlay (2018)) or GVOs (Brown et al., 2022). This would ensure that both the input fields and the inferred flows are independent of a global representation, providing a more stringent test of the extent to which the recovered flow structures are supported by regional observations alone. It would also allow the robustness of the methodology to observational noise, residual external-field contamination, and uncertainties in geomagnetic field estimation to be assessed more directly.

## 7.2 Do localised flow models contain information helpful for forecasting the Secular Variation of the Earth's magnetic field?

In Chapter 5, we developed regional and global core-flow models to produce a candidate secular-variation forecast for the 14th generation of the IGRF. This was done by producing regional and global core-flow models using the PINN approach detailed in Chapter 4. These flows were then used to advect the magnetic field under the assumption of steady flow, producing hindcasts of the secular variation over the IGRF-13 period. Although the global inversion reproduced the initial secular variation more closely, the regional model retained more small-scale flow structure. Despite these differences, the resulting regional and global secular-variation spectra were very similar, with relative differences of 1.8% at the Earth's surface and 1.7% at the CMB.

The regional approach produced the more accurate hindcast. When the predicted magnetic fields were compared with CHAOS-7.18, the regional model had a lower RMS difference than both the global model and the other IGRF-13 candidate models at the CMB. It also had the lowest RMS difference at the Earth's surface, whereas several IGRF-13 candidates outperformed the global model. The same relative performance was found when the hindcasts were compared with Kalmag rather than CHAOS-7.18, indicating that the result was not solely caused by the choice of reference field model. The improved hindcast produced by the regional approach therefore suggests that the additional small-scale flow structure retained by the local inversions contains useful information for forecasting secular variation.

Following the publication of the 14th generation of the IGRF (Beggan et al., 2026), the updated Leeds forecast can now be compared with the other IGRF-14 candidate models. The results presented here use the corrected Leeds candidate rather than the version submitted before

the deadline, which contained an error in the advection calculation. The Mauersberger–Lowes spectra of the candidate secular-variation forecasts are shown in Figure 7.1. The updated Leeds candidate is shown by the solid black line, while the mean and median candidate models are shown by the solid green and purple lines, respectively. The Leeds spectrum is close to both the mean and median, showing that the regional core-flow methodology produces a forecast that is consistent with the wider group of IGRF-14 candidates rather than being an outlier. This is further supported by Figure 7.2, which shows the difference between each candidate forecast and the median of all candidates. The spatial differences associated with the Leeds candidate are generally small compared with those of several other candidate models, placing it close to the median forecast.

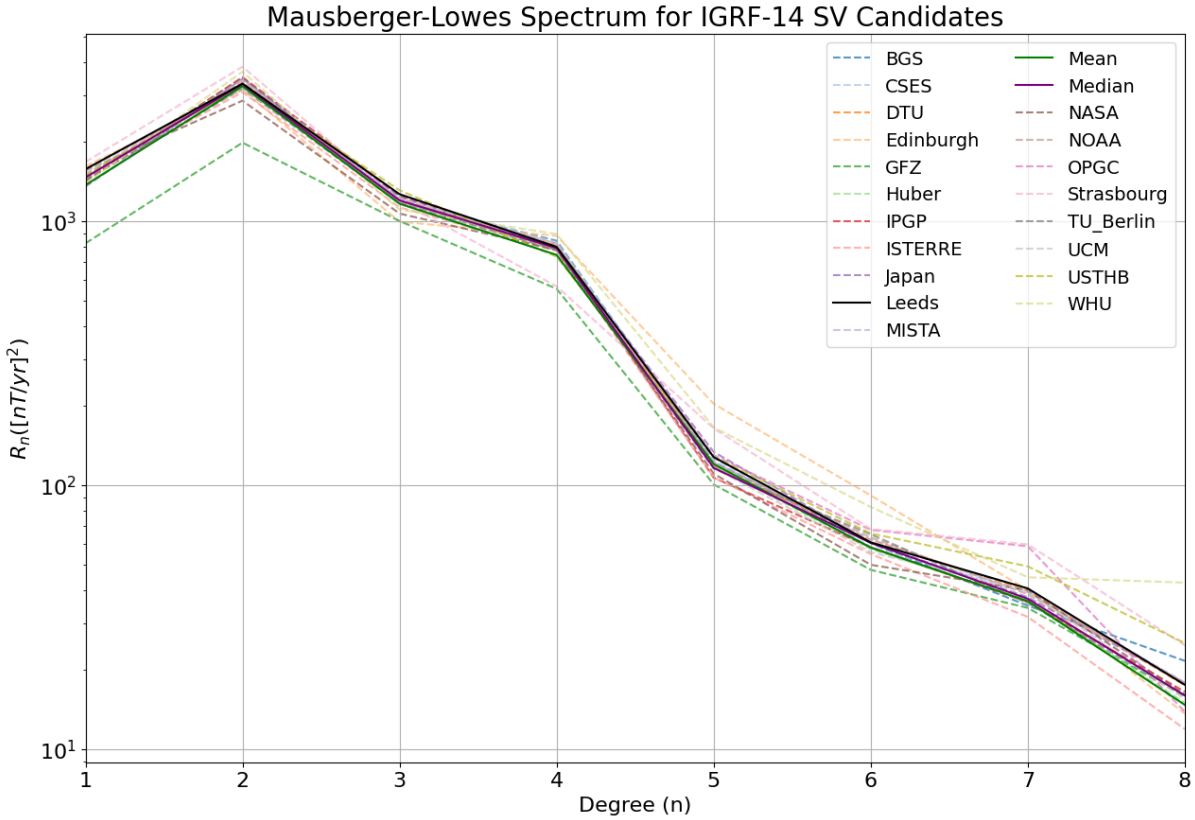



Figure 7.1: Mauersberger–Lowes spectra of the secular-variation forecast candidates for IGRF-14. The mean and median of the candidates are shown by the solid green and purple lines, respectively, while the updated Leeds candidate is shown by the solid black line. Figure created by adapting code from the IGRF-14 evaluation repository (IAGA V-MOD).

Overall, the improved hindcast produced by the regional model shows that localised flow inversions can retain information that is useful for forecasting the Earth's secular variation. Dividing the CMB into independently inverted regions allowed the model to preserve regional flow structure that was smoothed in the global inversion, while still producing a coherent global forecast.

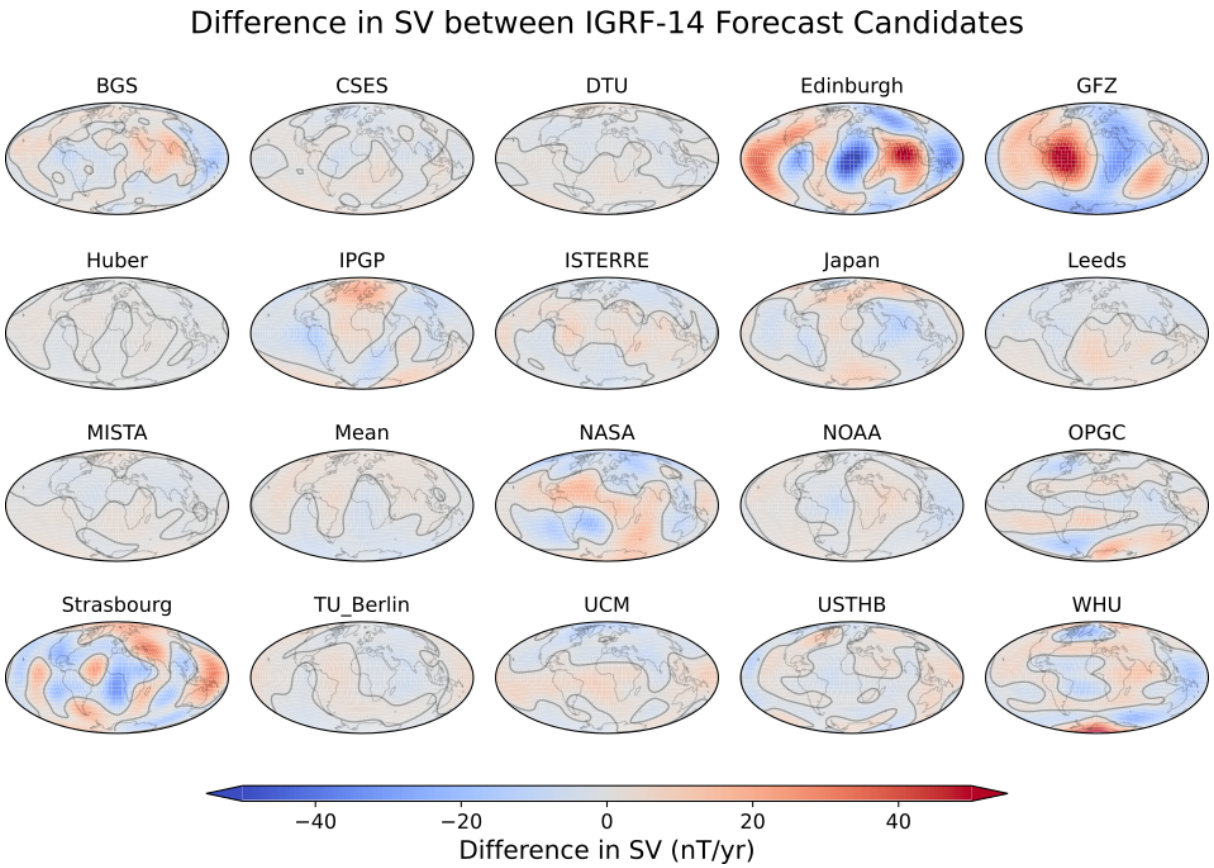



Figure 7.2: Difference between each IGRF-14 candidate radial secular-variation forecast and the median of all candidate forecasts. Figure created by adapting code from the IGRF-14 evaluation repository (IAGA V-MOD), plotted at the Earth's surface.

This improved hindcast performance motivated the use of the regional methodology to construct the University of Leeds candidate secular-variation model for IGRF-14 over 2025–2030. Its agreement with the mean and median IGRF-14 candidates further shows that the regional methodology produces a plausible forecast that is consistent with forecasts obtained using other approaches. However, whether the improved hindcast performance will persist over the IGRF-14 forecast period remains unknown. The difference between the regional model and the next best-performing candidate in the hindcast was relatively small, and unpredicted changes in secular variation, including those associated with geomagnetic jerks, could reduce the performance of the steady-flow advection approach. Further work is also required to determine whether the improved hindcast arises from the regional inversion itself or partly from differences between the grids used to train the regional and global models. Testing both approaches using equivalent grids and alternative geomagnetic field models would provide a stronger assessment of the forecasting information contributed specifically by the regional methodology.

## 7.3 Can heat flux heterogeneity be inferred from the radial magnetic field at the Core Mantle Boundary?

In Chapter 6, we developed a convolutional autoencoder classifier to test whether CMB heat-flux heterogeneity could be inferred from snapshots of the radial magnetic field. The methodology was first tested using the controlled LEDCJ simulations, before being applied to a wider set of simulations with varying dynamo parameters and heat-flux patterns, and finally to the CHAOS-8.5 and gufm1 models of the Earth's magnetic field. The controlled results showed that the classifier could distinguish between homogeneous, weakly heterogeneous and strongly heterogeneous cases with a validation error of 4.9%. Classification depended on relative field strength and non-dipolar magnetic structure, with performance decreasing below spherical harmonic degree $n = 4$. Regional tests and saliency maps showed that the most useful information was concentrated at equatorial and mid-latitudes, particularly near the Pacific and African LLVPs.

The validation error increased to 20.1% for the wider simulation set, but decreased to 12.9% when cases with $Q^* < 1$ were removed. This indicates that stronger heat-flux heterogeneity produces a clearer magnetic signature, while weakly heterogeneous cases may be difficult to distinguish from the natural variability of homogeneous simulations. Overall, the results show that the radial magnetic field contains classifiable information about imposed CMB heat-flux heterogeneity, but that this information becomes harder to isolate when the dynamo regime and heat-flux pattern also vary.

The application to the Earth's magnetic field showed that the classifications were not sufficiently robust to infer the structure of CMB heat flux. Although CHAOS-8.5 was consistently classified as corresponding to homogeneous CMB heat flux, the gufm1 predictions varied between epochs and random initialisations. The predicted class therefore depended on the geomagnetic field model used, rather than remaining consistent between different representations of the Earth's field. The saliency maps also showed that the CHAOS-8.5 classifications were strongly influenced by magnetic-field structure near the South Atlantic Anomaly, despite this feature not being well represented in the training simulations. This does not necessarily mean that the classification is incorrect, but it makes it difficult to determine whether the classifier has identified a physically meaningful heat-flux signature or is responding to a feature that it did not encounter during training. This demonstrates the importance of investigating the otherwise black-box nature

of machine-learning models, particularly in scientific applications where predictions must be connected to physically meaningful features.

Improving the robustness of the Earth classifications will require a better understanding of the information used by the classifier. Additional explainable AI methods should be applied alongside Grad-CAM, since different approaches test different aspects of the model and no single method provides a complete explanation (Barredo Arrieta et al., 2020; Skliarov et al., 2025). Occlusion tests could mask selected regions to determine whether removing features such as the South Atlantic Anomaly changes the predicted class. Integrated Gradients and SmoothGrad could test whether the regions highlighted by Grad-CAM remain important and are stable to small changes in the input field (Sundararajan et al., 2017; Smilkov et al., 2017). SHAP could also estimate whether broader, physically defined regions, such as latitude bands or areas surrounding the LLVPs, support a homogeneous or heterogeneous classification (Lundberg and Lee, 2017). Agreement between these methods would provide greater confidence that the classifier is responding to physically meaningful magnetic-field structure, whereas disagreement would indicate that the interpretation remains dependent on the explanation method used.

The simulation-to-reality gap must also be reduced by expanding the training dataset. The classifier was trained using pre-existing simulations that were produced for different scientific purposes, rather than a controlled ensemble designed specifically for application to the Earth. A more suitable ensemble would systematically vary the dynamo regime, the amplitude and spatial pattern of CMB heat flux, and the morphology and spectral distribution of the resulting magnetic field. It should include simulations spanning homogeneous to strongly heterogeneous heat flux, with other parameters held fixed where possible, as well as simulations in which individual control parameters are varied independently. This would make it easier to separate magnetic-field changes caused by heat-flux heterogeneity from those caused by the broader dynamo regime. Training across a deliberately broad range of simulated conditions would also increase the likelihood that the Earth's magnetic field falls within the range represented during training (Tobin et al., 2017; Chen et al., 2022).

Producing such a training ensemble would, however, be computationally demanding. A systematic study would require several heat-flux amplitudes and spatial patterns to be tested across multiple dynamo regimes, with each simulation integrated for long enough to produce representative magnetic-field snapshots. The appropriate range of simulations is also uncertain because

reproducing an Earth-like magnetic-field morphology does not guarantee that a simulation has the same underlying force balance as the Earth's core (Nakagawa and Davies, 2022). Rather than attempting to represent every possible parameter combination, a more targeted approach could identify the largest differences between the existing simulations and CHAOS-8.5 or gufm1, and prioritise new simulations that reduce these differences. Additionally, the training simulations were not selected using other measures of Earth-like field behaviour, such as $f_{\mathrm{dip}}$, and alternative selection criteria may therefore influence the results.

Overall, the results suggest that CMB heat-flux heterogeneity can be inferred from the radial magnetic field when its magnetic signature is sufficiently strong and represented within the simulation-based training data. The controlled and wider simulation results provide evidence that heat-flux heterogeneity leaves classifiable information in the morphology and amplitude of the radial magnetic field. However, the lack of agreement between the CHAOS-8.5 and gufm1 classifications means that this conclusion cannot yet be extended reliably to the Earth. Further work is therefore needed to broaden the simulation ensemble, reduce the simulation-to-reality gap and establish that the classifier is using physically meaningful magnetic-field features before it can provide a robust inference of the Earth's CMB heat-flux structure.

## 7.4 Concluding Remarks and Further Work

This thesis has demonstrated how machine-learning methods can be developed and applied to extract geophysical information from the Earth's magnetic field, providing new approaches for investigating processes at the top of the core and the core–mantle boundary. The novel contributions lie in the development of machine-learning methods for three related geomagnetic problems. Physics-informed neural networks were used to perform regional core-flow inversions from locally restricted magnetic-field information, providing an independent test of structures previously identified using global methods. These regional flows were then used to forecast secular variation, showing that localised flow models can retain information that is useful for prediction. Finally, a convolutional autoencoder classifier was developed to identify signatures of imposed CMB heat-flux heterogeneity in simulated radial magnetic fields. Together, these approaches demonstrate that machine learning can provide flexible solutions to a range of geomagnetic problems, while allowing physical constraints, spatial information and model interpretation to be incorporated into the analysis.

The results also demonstrate that the use of machine learning in geomagnetism requires caution. The flexibility that allows these methods to identify complex relationships can also make it difficult to determine why a particular solution or prediction has been obtained. This was especially clear when the heat-flux classifier was applied to geomagnetic field models, where the results depended on the model and random initialisation used, and where the saliency maps highlighted features that were not well represented in the training simulations. Machine-learning results should therefore not be interpreted from their predictive performance alone. Their robustness must be tested across data, initialisations and modelling choices, and the features used by the model must be examined to determine whether they have a physically meaningful interpretation. Addressing this black-box behaviour will be important if machine-learning methods are to become trusted and widely accepted within the geomagnetic community.

Future progress in this field would benefit from shared benchmarks that allow machine-learning methods to be compared using common geomagnetic problems, datasets and evaluation measures. These should include synthetic tests where the correct solution is known, hindcasting experiments, sensitivity to random initialisation, and tests using more than one geomagnetic field model. Explainable AI methods should also become a standard part of scientific machine-learning studies, with several complementary approaches used to test whether the same physical features are consistently identified. Finally, publishing numerical simulations, processed training datasets, trained models and analysis code under the FAIR (Findable, Accessible, Interoperable and Reusable) principles would enable results to be reproduced and larger, more representative training ensembles to be constructed (Wilkinson et al., 2016). This approach is already being adopted in geomagnetism: ISTerre has developed a collaborative web tool that catalogues existing dynamo simulations (Claveau, 2025), while Rogers et al. (2026) published the full set of simulations used in their study. Collectively, the adoption of shared benchmarks, explainable AI and FAIR research practices would support more robust, reproducible and physically interpretable applications of machine learning in geomagnetism.

## References

Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F., 2020. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion 58, 82–115. doi:10.1016/j.inffus.2019.12.012.

Beggan, C.D., Kloss, C., Amblard, P., Amezza, T., Amit, H., Aubert, J., Baerenzung, J., Bregnhøj Lauridsen, J., Brown, W., Califf, S., Campuzano, S.A., Chauvet, L., Chulliat, A., Claveau, R., Coïsson, P., da Silva, M., Davies, C., Deborde, R., Dobrovolsky, M., Fillion, M., Finlay, C.C., Firsov, I., Fournier, A., Gailler, L., Gillet, N., Gomez Perez, N., Grayver, A., Gwirtz, K., Hamoudi, M., Holme, R., Holschneider, M., Hulot, G., Jager, T., Jiang, Y., Kada, M., Korte, M., Krasnoperov, R., Kuang, W., Kudin, D., Léger, J.M., Lemgharb, A., Lesur, V., Liu, P., Livermore, P.W., Madsen, F.D., Mandea, M., Matsushima, M., Matzka, J., Michaelis, I., Minami, T., Nair, M., Nakano, S., Oehler, J.F., Olsen, N., Pavón-Carrasco, F.J., Presnyakov, S., Qian, B., Rang, X., Rogers, H.F., Rother, M., Rouxel, D., Sabaka, T., Sato, S., Saynisch-Wagner, J., Schanner, M., Serrano, M., Shakespeare-Rees, N., Shi, Y., Smith, A., Soloviev, A., Terra-Nova, F., Thébault, E., Tøffner-Clausen, L., Toh, H., Wang, J., Wardinski, I., Watson, C., Whaler, K., Xiong, C., Yan, Q., Yang, Y., Yao, H., Zeren, Z., Zhang, K., Zhou, B., 2026. International geomagnetic reference field: The fourteenth generation. Earth, Planets and Space 78, 127. URL: https://doi.org/10.1186/s40623-025-02360-0, doi:10.1186/s40623-025-02360-0.

Berrone, S., Canuto, C., Pintore, M., Sukumar, N., 2023. Enforcing Dirichlet boundary conditions in physics-informed neural networks and variational physics-informed neural networks. Heliyon 9, e18820. doi:10.1016/j.heliyon.2023.e18820.

Bischof, R., Kraus, M.A., 2025. Multi-objective loss balancing for physics-informed deep learning. Computer Methods in Applied Mechanics and Engineering 439, 117914. doi:10.1016/j.cma.2025.117914.

Brown, W., Beggan, C., Gomez Perez, N., Hammer, M., Finlay, C., Hulot, G., 2022. Geomagnetic virtual observatories for investigating sub-annual core field variation. URL: https://lps22.esa.int/frontend/index.php?page_id=18477.

Buffett, B., Knezek, N., Holme, R., 2016. Evidence for MAC waves at the top of Earth's core and implications for variations in length of day. Geophysical Journal International 204, 1789–1800. URL: https://doi.org/10.1093/gji/ggv552, doi:10.1093/gji/ggv552, arXiv:https://academic.oup.com/gji/article-pdf/204/3/1789/1988933/ggv552.pdf.

Chen, X., Hu, J., Jin, C., Li, L., Wang, L., 2022. Understanding domain randomization for sim-to-real transfer, in: International Conference on Learning Representations. URL: https://openreview.net/forum?id=T8vZHIRTrY.

Claveau, R., 2025. Planetary dynamo simulation explorer. URL: https://gricad-gitlab.univ-grenoble-alpes.fr/claveaur/planetary-dynamo-simulation-explorer. gitLab repository, accessed 12 August 2026.

Finlay, C.C., Gillet, N., Aubert, J., Livermore, P.W., Jault, D., 2023. Gyres, jets and waves in the earth's core. Nature Reviews Earth & Environment 4, 377–392. URL: https://www.nature.com/articles/s43017-023-00425-w, doi:10.1038/s43017-023-00425-w.

Gillet, Gerick, F., Jault, D., Schwaiger, T., Aubert, J., Istas, M., 2022. Satellite magnetic data reveal interannual waves in earth's core. Proceedings of the National Academy of Sciences of the United States of America 119, e2115258119. doi:10.1073/pnas.2115258119.

Grüne, C.R., Whaler, K.A., Madsen, F.D., 2025. Detecting low-latitude outer core-surface waves with 25 years of satellite secular variation data. Physics of the Earth and Planetary Interiors 367, 107435. URL: https://www.sciencedirect.com/science/article/

pii/S0031920125001293, doi:https://doi.org/10.1016/j.pepi.2025.107435.

Hammer, M.D., Finlay, C.C., 2018. Local averages of the core–mantle boundary magnetic field from satellite observations. Geophysical Journal International 216, 1901–1918. URL: https://doi.org/10.1093/gji/ggy515, doi:10.1093/gji/ggy515, arXiv:https://academic.oup.com/gji/article-pdf/216/3/1901/27493010/ggy515.pdf.

Horne, M.J.S., Jimack, P.K., Khan, A., Wang, H., 2025. Hard constraint projection in a physics informed neural network, in: Proceedings of the 35th International Conference on Parallel Computational Fluid Dynamics, Forschungszentrum Jülich GmbH, Zentralbibliothek. pp. 50–53. doi:10.34734/FZJ-2025-02453.

IAGA V-MOD, . IGRF14eval: Evaluation files for IGRF-14. GitHub repository. URL: https://github.com/IAGA-VMOD/IGRF14eval/tree/main. accessed 10 August 2026.

Kloss, Finlay, 2019. Time-dependent low-latitude core flow and geomagnetic field acceleration pulses. Geophysical Journal International 217, 140–168. doi:10.1093/gji/ggy545.

Lagaris, I.E., Likas, A., Fotiadis, D.I., 1998. Artificial neural networks for solving ordinary and partial differential equations. IEEE Transactions on Neural Networks 9, 987–1000. doi:10.1109/72.712178.

Lundberg, S.M., Lee, S.I., 2017. A unified approach to interpreting model predictions, in: Advances in Neural Information Processing Systems.

Luong, K.A., Le-Duc, T., Lee, J., 2024. Automatically imposing boundary conditions for boundary value problems by unified physics-informed neural network. Engineering with Computers 40, 1717–1739. doi:10.1007/s00366-023-01871-2.

McClenny, L.D., Braga-Neto, U.M., 2023. Self-adaptive physics-informed neural networks. Journal of Computational Physics 474, 111722. doi:10.1016/j.jcp.2022.111722.

Nakagawa, T., Davies, C.J., 2022. Combined dynamical and morphological characterisation of geodynamo simulations. Earth and Planetary Science Letters 594, 117752. doi:10.1016/j.epsl.2022.117752.

Pais, M.A., Jault, D., 2008. Quasi-geostrophic flows responsible for the secular variation of the earth's magnetic field. Geophysical Journal International 173, 421–443. URL: https://doi.org/10.1111/j.1365-246X.2008.03741.x, doi:10.1111/j.1365-246X.2008.03741.x, arXiv:https://academic.oup.com/gji/article-pdf/173/2/421/6078007/173-2-421.pdf.

Rogers, H., Davies, C., Constable, C., Mason, S., Frasson, T., 2026. Using paleomagnetic observations and geodynamo simulations to assess earth's magnetic field morphology and variability. Journal of Studies of the Earth's Deep Interior 2. URL: https://doi.org/10.5281/zenodo.21821109, doi:10.5281/zenodo.21821109.

Shakespeare-Rees, N., Livermore, P.W., Davies, C.J., Rogers, H.F., Beggan, C.D., Brown, W.J., Finlay, C.C., 2026. Forecasting secular variation using physics-informed neural networks for IGRF-14. Earth, Planets and Space 78, 98. URL: https://doi.org/10.1186/s40623-026-02427-6, doi:10.1186/s40623-026-02427-6.

Shakespeare-Rees, N., Livermore, P.W., Davies, C.J., Rogers, H.F., Brown, W.J., Beggan, C.D., Finlay, C.C., 2025. Local flow estimation at the top of the earth's core using physics informed neural networks. Physics of the Earth and Planetary Interiors 367, 107424. URL: https://www.sciencedirect.com/science/article/pii/S0031920125001189, doi:https://doi.org/10.1016/j.pepi.2025.107424.

Sheikholeslami, M., Salehi, S., Mao, W., Eslamdoost, A., Nilsson, H., 2025. Physics-informed neural networks with hard and soft boundary conditions for linear free surface waves. Physics of Fluids 37, 087158. doi:10.1063/5.0277421.

Skliarov, M., El Shawi, R., Dhaoui, C., Ahmed, N., 2025. A comparative evaluation of explainability techniques for image data. Scientific Reports 15, 41898. doi:10.1038/

s41598-025-25839-y.

Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M., 2017. SmoothGrad: Removing noise by adding noise. arXiv preprint arXiv:1706.03825 doi:10.48550/arXiv.1706.03825.

Sukumar, N., Srivastava, A., 2022. Exact imposition of boundary conditions with distance functions in physics-informed deep neural networks. Computer Methods in Applied Mechanics and Engineering 389, 114333. doi:10.1016/j.cma.2021.114333.

Sundararajan, M., Taly, A., Yan, Q., 2017. Axiomatic attribution for deep networks, in: Proceedings of the 34th International Conference on Machine Learning, PMLR. pp. 3319–3328.

Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P., 2017. Domain randomization for transferring deep neural networks from simulation to the real world, in: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE. pp. 23–30. doi:10.1109/IROS.2017.8202133.

Wang, S., Teng, Y., Perdikaris, P., 2021. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. SIAM Journal on Scientific Computing 43, A3055–A3081. doi:10.1137/20M1318043.

Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.W., da Silva Santos, L.B., Bourne, P.E., Bouwman, J., Brookes, A.J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C.T., Finkers, R., Gonzalez-Beltran, A., Gray, A.J.G., Groth, P., Goble, C., Grethe, J.S., Heringa, J., 't Hoen, Peter A. C., Hooft, R., Kuhn, T., Kok, R., Kok, J., Lusher, S.J., Martone, M.E., Mons, A., Packer, A.L., Persson, B., Rocca-Serra, P., Roos, M., van Schaik, R., Sansone, S.A., Schultes, E., Sengstag, T., Slater, T., Strawn, G., Swertz, M.A., Thompson, M., van der Lei, J., van Mulligen, E., Velterop, J., Waagmeester, A., Wittenburg, P., Wolstencroft, K., Zhao, J., Mons, B., 2016. The FAIR guiding principles for scientific data management and stewardship. Scientific Data 3, 160018. URL: https://doi.org/10.1038/sdata.2016.18, doi:10.1038/sdata.2016.18.

# Appendix A

# Appendix Chapter

## A.1 Toroidal Flows

While the results in Chapter 4 used the Tangentially Geostrophic flow assumption, other flow assumptions can be used within the PINN framework. To implement the toroidal-only flow constraint, the poloidal scalar ($P$) is set to zero, such that the horizontal flow components can be written as

$$\boldsymbol{u}_\theta = \frac{1}{\sin\theta}\frac{\partial T}{\partial \phi}, \tag{A.1}$$

$$u_\phi = -\frac{\partial T}{\partial \theta}. \tag{A.2}$$

The network is therefore configured to calculate only the Toroidal scalar $T$, with the sole loss term described in Equation(4.11). In this formulation, the toroidal-only flow assumption is encoded as a hard constraint. Alternatively, it could be imposed as a soft constraint by allowing the network to fit both $T$ and $P$ while including a loss term that encourages $P$ to approach zero.

Figure A.1 shows the resulting flows, SV, and SV residuals for the hard constraint approach, using CHAOS-8.1 in January 2024. These results were determined using the same network size as the models described in Chapter 4, and so further work could be to find the optimum network size for this flow assumption. The MAE between the SV from the CHAOS-8.1 model and the SV output by the PINN is 396nT/year (Relative L1 Error of 10.6%).

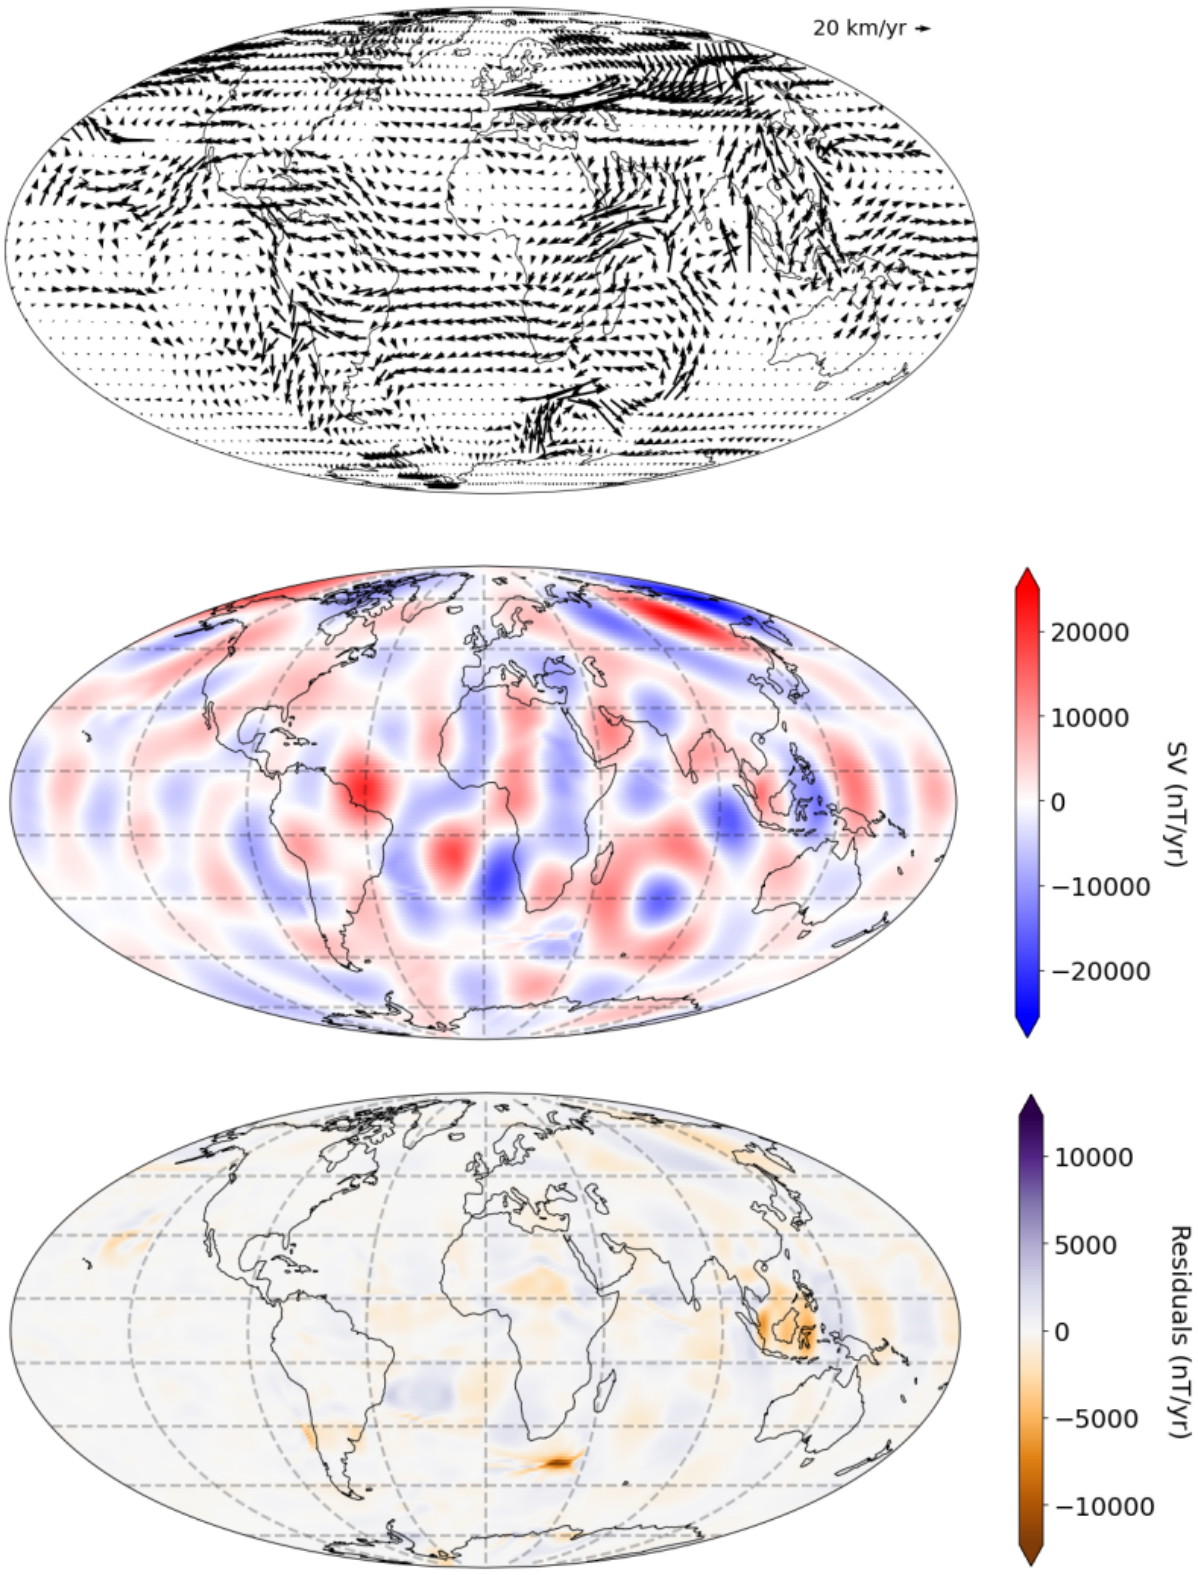



Figure A.1: (Top) Global patchwork of toroidal regional flows in January 2020. The thickness of the arrows indicate flow magnitude, and the arrows indicate direction. (Middle) SV calculated using the recovered flows in the top panel. (Bottom) Residuals between the SV from CHAOS-8.1 and the SV shown in the middle panel. The box boundaries (after the 5◦ border removal) are marked by grey dashed lines.

## A.2 TG Misfit

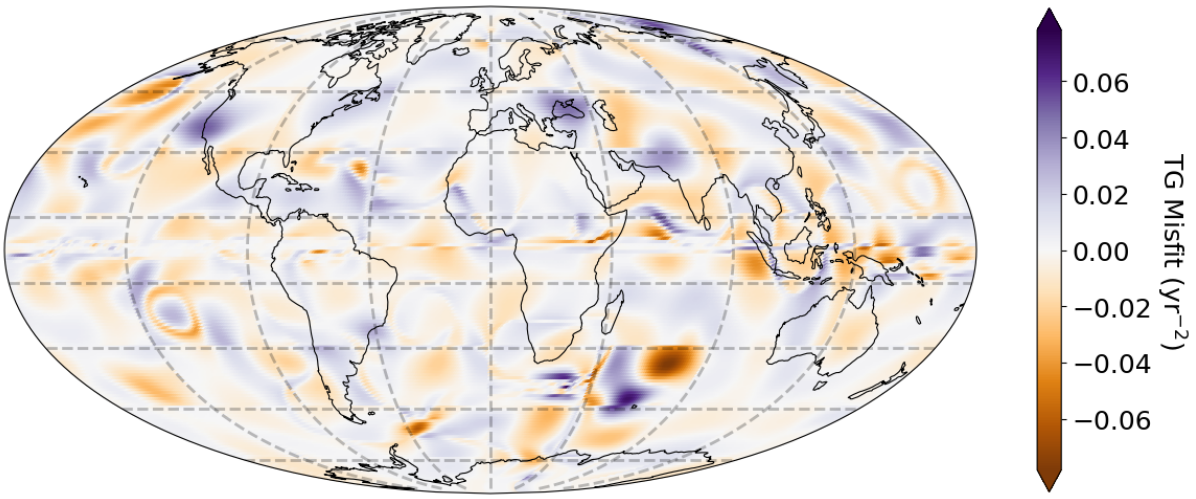



Figure A.2: TG misfit for the global mosiac shown in Figure 4.9, in squared units as the misfit is quantified by the MSE.