'Advances in machine learning and data analysis methods for solar, solar wind, and space weather applications with an emphasis on physical insight from large and complex datasets (MLSolar)' Abstracts

Talks

Characterising Pre-Eruptive Magnetic Conditions in the Lower Solar Atmosphere with ES-MAGNAR

Authors

Marianna Korsos (1)

(1) University of Sheffield

Abstract

Solar active regions are the primary sources of solar flares and coronal mass ejections; however, only a small fraction of magnetically complex regions produce major eruptions. Identifying the magnetic conditions that distinguish eruptive from non-eruptive active regions remains a key challenge for space-weather forecasting.

Recently, we found that extending the application of solar flare predictor parameters from the photosphere into the lower solar atmosphere improves the identification of eruptive active regions. Building on this result, we introduce our newly developed machine-learning-based ES-MAGNAR (Eruptive State – MAGnetic Nonlinear AutoRegressive) prediction model, based on a NARX-type (Nonlinear AutoRegressive model with eXogenous inputs) architecture.

The model incorporates multi-height magnetic diagnostics and predicts the eruptive state of active regions 3, 6, 12, and 24 hours in advance. The predictions are based on the temporal evolution of ten magnetic morphological parameters, evaluated at multiple atmospheric heights between the photosphere and the lower corona (~7 Mm). Our results show that the ES-MAGNAR model can reliably predict the eruptive state of active regions associated with major solar eruptions.

A machine learning framework for joint Hanle–Zeeman spectropolarimetric inversions

Authors

Ryan Campbell (1), tanausú del pino alemán (2), javier trujillo bueno (2)

(1) Queen's University Belfast
(2) instituto de astrofisica de canarias

Abstract

The Hanle effect is often invoked as a complementary diagnostic to the Zeeman effect for probing weak and unresolved magnetic fields in the solar atmosphere. However, its practical utility for quantitative inversions remains uncertain due to degeneracies and the inherently non-local nature of line formation. We present a controlled machine-learning experiment designed to test whether Hanle-sensitive spectral lines provide additional constraints on atmospheric parameters beyond standard Zeeman diagnostics. A Transformer-based inversion model is trained on synthetic spectropolarimetric data to recover stratified atmospheric parameters from three configurations: (i) Zeeman-sensitive lines alone (Fe I 525 nm), (ii) Hanle-sensitive lines alone (Sr I 460 nm), and (iii) their combination. By comparing performance across identical training and evaluation conditions, we isolate the contribution of Hanle diagnostics to inversion accuracy and uncertainty. Preliminary results will be presented evaluating whether the inclusion of Hanle-sensitive lines leads to measurable improvements in magnetic and thermodynamic parameter inference. Particular attention is given to how these conclusions depend on modelling assumptions, including radiative transfer dimensionality and the treatment of spatial context. This work provides a framework for using machine-learning inversions not only as predictive tools, but also as controlled experiments to quantify the information content of spectropolarimetric observables.

Forecasting of Photospheric Magnetograms using deep learning with Farside seismic priors

Authors

Hemapriya Raju (1), Dattaraj B Dhuri (2), Shravan Hanasoge (3)

(1) Tata Institute of Fundamental Research
(2) Center for Astrophysics and Space Science, New York University Abu Dhabi
(3) Center for Astrophysics and Space Science, New York University Abu Dhabi, Department of Astronomy and Astrophysics, Tata Institute of Fundamental Research,

Abstract

Accurate knowledge of the global photospheric magnetic field is essential for modelling the solar
corona and forecasting space weather. As we are limited by the nearside observations from instruments such as the Helioseismic and Magnetic Imager (HMI) onboard the Solar Dynamics Observatory, getting a reliable global magnetic field is challenging. Integrating farside observations have improved that scenario. In this work, we propose a UNet based deep learning model, that takes one solar rotation of magnetograms as inputs, and farside maps as auxiliary input. The model's objective is to predict future photospheric magnetograms for next few days initially, so that in long run, it can be seen as an alternative to replace the farside maps. Temporal evolution is incorporated by treating sequential magnetograms as additional input channels, allowing the network to capture the spatiotemporal evolution of the magnetic field. Since far-side seismic maps are prone to noise, threshold-based saturation is applied to extract reliable information about large magnetic regions. The current model could reproduce large scale transport processes such as differential rotation and could well reproduce the polarity of magnetic structures. In few cases, where farside active region is more stable, the model is able to capture the emergence of large active regions rotating onto the visible disk. This model can aid in forecasts of solar magnetic field evolution, and provide longer lead time for space weather predictions.

Predicting Solar Flares with a Magnetic Topology ML Model

Authors

Thomas Williams (1), Christopher Prior (1), David MacTaggart (2), Aslam O.P.M. (2), D. Shaun Bloomfield (3)

(1) Durham University
(2) University of Glasgow
(3) Northumbria University

Abstract

Solar flares are a primary driver of space weather, and forecasting their occurrence remains a significant challenge. This work presents a novel flare prediction model based on topologically derived photospheric magnetic parameters. We employ the \texttt{ARTop} framework to compute the time-dependent input rates of magnetic winding and helicity across $>10^5$ active region observations, decomposing them into current-carrying and potential components to reduce sensitivity to optical flow methods. An \texttt{XGBoost} machine learning model is trained on these time series, alongside engineered features including rolling statistics, kurtosis, and flare history, to predict the probability of $\ge$M1.0 class flares within the next 24 hr. The model demonstrates strong performance on a validation set, with a true skill statistic (TSS) of 0.804 for once-daily operational region forecasts. When applied to a fully independent dataset, the model achieves a TSS of 0.524. A SHAP analysis confirms the model's physical interpretability, identifying flare history and accumulated current-carrying winding and helicity as the most important features. The main challenges identified are false positives arising from active regions with frequent C-class flaring and systematic errors via projection effects when active regions are near the limb. Excluding limb-affected data yields no improvement in the holdout set TSS (0.521 versus 0.524), due to the overall decreased number of flares. However, our per-region analysis indicates that mitigating these projection effects is crucial for future operational deployment. This work establishes magnetic topology, particularly its current-carrying components, as a highly effective and physically meaningful set of predictors for solar flare forecasting.

Advancing Solar Cycle Forecasting with Novel Physics-Informed and Peak-Aware Transformer Models

Authors

Sherwin D Souza (1), Dr. Amrita Prasad

(1) Cardiff Metropolitan University

Abstract

Solar activity varies across a wide range of temporal scales, from short-term fluctuations to long-term cyclic behaviour. The 13-month smoothed sunspot number is commonly used in solar cycle studies to filter out high-frequency variability and highlight the underlying cycle structure, making it a standard reference for both analysis and prediction. However, this smoothing process removes short-term variations that can be significant, particularly near solar maxima. In contrast, monthly mean sunspot numbers retain these rapid changes, which may reflect important physical processes associated with solar dynamics. For this reason, both smoothed and monthly datasets are considered in this study. The smoothed series supports the learning of long-term cycle behaviour, while the monthly data provides a more realistic and challenging setting for evaluating predictive performance under higher variability.

The approaches include a transformer baseline model, a hybrid physics-informed and a novel peak-aware transformer model. Both approaches incorporate a physics-based smoothness constraint during training, capturing the known local continuity of solar cycle evolution and embedding it as an inductive bias. The peak-aware formulation further introduces an asymmetric boundary loss designed to mitigate the "peak-shaving" effect and improve the representation of solar maxima and minima. A rigorous walk-forward evaluation strategy is employed, in which each model is trained exclusively on historical cycles and evaluated on the subsequent unseen cycle, ensuring that all reported metrics reflect true out-of-sample performance. Hyperparameter optimisation is conducted independently for each approach, and performance is assessed using RMSE and MAE across individual solar cycles.

While the proposed Hybrid Physics Transformer Model performs strongly on smoothed sunspot data, the increased variability in monthly mean sunspot numbers leads to higher RMSE values. Our evaluation across Solar Cycles 22–25 reveals a distinct trade-off: while the physics-informed Transformer achieves a lower overall RMSE by maintaining a smoother trajectory, it consistently fails to reach maximum solar amplitudes. In contrast, the novel peak-aware transformer model demonstrates significantly higher functional utility by more effectively capturing solar maximum amplitudes present in the monthly mean sunspot data. By prioritizing cycle extremes during the training process, this model successfully mitigates the characteristic under-prediction of solar peaks. Nevertheless, the models maintain improved predictive capability relative to existing approaches, particularly in capturing cycle amplitude and extreme events. This demonstrates the robustness of the proposed framework across data representations and highlights its potential for both scientific analysis and operational forecasting.

Constraining CME model parameters with data assimilation of heliospheric imager observations

Authors

Matthew Lang (1), Luke Barnard (1)

(1) University of Reading

Abstract

Space weather describes the variable conditions in near Earth space and their impacts on humans and our technologies. Among the most important drivers of space weather impacts are coronal mass ejections (CMEs), which can generate severe geomagnetic storms, induce currents in the Earth's magnetosphere and upper atmosphere, and cause damage to satellites, communication infrastructure, and power grids. Accurate forecasting of CMEs is therefore essential for effective mitigation of these hazards.

Despite significant advances in heliospheric modelling, numerical simulations of CME evolution remain highly uncertain, with typical arrival time errors of the order of 10 hours. A major source of this uncertainty arises from the CME parameterisations within solar wind models, where input parameters are inferred from remote observations made within ~0.1 AU of the Sun. These parameters therefore only provide limited constraints on CME evolution during their propagation into the solar system. Further compounding this uncertainty is the scarcity of in situ measurements upstream of Earth. At present, consistent upstream monitoring is provided almost exclusively by spacecraft at L1, offering only 1–2 hours of warning prior to CME impact which is insufficient for many operational responses. To reduce these uncertainties, it is essential to make greater use of remote observations spanning a greater proportion of the heliosphere. In this work, we present developments to the SIR HUXt particle filter data assimilation framework, incorporating CME elongation profiles derived from Heliospheric Imager data to constrain CMEs as they propagate towards Earth. We demonstrate the performance of this approach using experiments based on real CME elongation observations and show how assimilation of heliospheric imager data can improve CME forecasts at Earth.

Posters

Connecting hybrid plasma simulations of collisionless shockwaves to in situ observations with machine learning

Authors

Imogen Gingell (1)

(1) University of Southampton

Abstract

Both observations and simulations have revealed that magnetic reconnection occurs at thin current sheets within the transition region of collisionless shock waves. These ion- and electron-scale structures arise from stream instabilities and turbulence in the shock layer, contribute significantly to repartition of energy across the shock, and propagate far downstream into the magnetosheath. In a recent study [Gingell et al. 2023, Physics of Plasmas, 30, 0123902], we used a series of 2D hybrid particle-in-cell simulations to explore the shock-driven generation and decay of reconnecting structures over a broad range of parameters. Magnetic field line integration was used to quantify reconnection in each simulation, classifying each cell in the domain as having “closed” or “open” magnetic field topology. Here, we use these classifications to train a convolution neural network (CNN) to identify regions of the simulation that are undergoing (or have undergone) magnetic reconnection. With a view to creating a dataset equivalent to a series of in situ observations, this is performed by splitting each simulation domain into a series of 1D virtual trajectories. We find that the trained CNN is able to effectively identify structures of interest in simulations that have different plasma and shock parameters to the training data set, as well as in those with different dimensionality (i.e. 3D simulations). Further, we present a pipeline for applying this simulation-trained CNN to in situ observations of shocks by the Magnetospheric Multiscale and Solar Orbiter spacecraft, and demonstrate successful detection of reconnection sites embedded in the shock layer. We discuss these techniques more generally as a case study for using machine learning to identify structures of interest in spacecraft data, which may contribute to on-board event selection for burst modes in spacecraft with relatively limited downlink capacity.

Turbulence and Intermittency within Interplanetary Coronal Mass Ejections

Authors

Himani Mehta (1), Zubair Shaikh (2), Shakti Singh (1), V S Pandey (1), Preeti Verma (1)

(1) National Institute of Technology Delhi, New Delhi, India
(2) William B. Hanson Center for Space Sciences, University of Texas at Dallas, Richardson, TX, USA

Abstract

Turbulence plays a central role in controlling how energy is transferred, dissipated, and ultimately converted into heat in space plasmas. Interplanetary Coronal Mass Ejections (ICMEs) provide a unique environment to study these processes, as they carry large-scale magnetic structures and generate a compressed sheath region via interaction with the ambient solar wind. In this work, we investigate the magnetic field turbulence and intermittency within and around several ICMEs observed between 2021-2025 by the Solar Orbiter mission till 1 AU. Using high-resolution magnetic field data, we examine the statistical properties of fluctuations in four distinct regions: the upstream solar wind, sheath, magnetic cloud, and trailing solar wind. Probability density functions (PDFs) are used to characterize the energy cascade, while structure functions and flatness quantify scale-dependent intermittency. We observed inertial-range spectral indices of approximately -1.64 in these regions, consistent with magnetohydrodynamic approximations, while the kinetic-scale spectra consistently steepen to about -3, indicating strong dissipation. Furthermore, fluctuations in the inertial range remain nearly Gaussian and multifractal, whereas those at kinetic scales become strongly non-Gaussian and monofractal. Comparisons across multiple events reveal that ICMEs tend to exhibit more stable and well-developed turbulence, while the surrounding sheath regions display steeper and more variable spectra, indicating ongoing energy transfer and turbulence evolution. This research deepens our understanding of how energy is distributed and dissipated in space plasmas, offering new perspectives for understanding processes that heat the solar wind, accelerate particles, and shape the broader behavior of space weather in our heliosphere.