Imperial College London

Empirical Investigation of Global Wildfire Drivers and Development of a New Flammability Parametrisation for the INFERNO Fire Model

---

Alexander Horst Ferdinand Kuhn-Régnier

February 2023

Imperial College London

Department of Physics

Submitted in part fulfilment of the requirements for the degree of

Doctor of Philosophy (PhD) of Imperial College London

and the Diploma of Imperial College London

Statement of Originality

All work presented herein is my own, with the exception of other work that is referenced as appropriate. The work presented has been published in part in the following article:

Alexander Kuhn-Régnier, Apostolos Voulgarakis, Peer Nowack, Matthias Forkel, I. Colin Prentice, and Sandy P. Harrison. The importance of antecedent vegetation and drought conditions as global drivers of burnt area. *Biogeosciences*, 18(12):3861–79. June 2021. ISSN 1726-4170. doi: 10.5194/bg-18-3861-2021.

Copyright Statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Wildfires have a significant impact on the Earth's vegetation, atmospheric composition, and climate. There is also growing evidence that fire behaviour has already been altered in response to climate change. Given the anticipated increase in climate conditions conducive to wildfires in many regions of the world, there is an urgent need to advance our understanding of the drivers of wildfires and improve their representation in global Earth system models. Recognising and predicting such responses to climate change and the associated feedbacks is one of the key challenges of the field, and integrated vegetation–fire models have the potential to accomplish this goal. However, the relationship between wildfire activity and past vegetation productivity is still unclear. It has been hypothesised that this strongly contributes to the poor performance of state-of-the-art fire models in representing observed vegetation–fire relationships. Thus, this thesis examines the role of seasonal and long-term vegetation dynamics (and fuel accumulation dynamics) that are of fundamental importance for global wildfires due to their impact on fuel availability during the fire season. In addition to instantaneous climatic conditions, the seasonality of antecedent vegetation and climate conditions controlling fuel build-up and fuel drying was found to be important for the prediction of burnt area in an empirical analysis of wildfire drivers. These were then used to modify the vegetation parametrisation of the INFERNO fire model. By using sigmoidal relationships to explicitly consider antecedent vegetation and climate conditions, global performance was improved. Finally, the new model was evaluated using sensitivity analyses, demonstrating the overarching importance of dryness and temperature while also disentangling the different impacts model parameters have on the magnitude and phase performance of the new parametrisation.

Acknowledgements

I would like to thank my supervisors Apostolos Voulgarakis, Sandy Harrison, and Colin Prentice for their continued support and guidance throughout the last four years, allowing me to pursue projects of previously inconceivable magnitude.

I also have to thank my parents, partner, and friends for their advice and steadfast encouragement throughout.

Publishing my first paper would not have been possible without the help of Peer Nowack and Matthias Forkel, who were a major source of inspiration. Additional thanks go to Kees Klein Goldewijk for contributing an updated HYDE dataset.

I am grateful to my examiners Heather Graven and Stephen Sitch for a stimulating discussion and advice leading to the latest version of this thesis.

I am thankful for the financial support from the NERC Centre for Doctoral Training in Quantitative and Modelling Skills in Ecology and Evolution (QMEE).

Code availability

Computer code was integral to the generation and visualisation of the results presented in Chapters 3–5 and can be found in the following `Python` packages:

- `wildfires` [Kuhn-Régnier, 2022c].
- `era5analysis` [Kuhn-Régnier, 2020].
- Chapter 3
  - `empirical-fire-modelling` [Kuhn-Régnier, 2021].
  - Accumulated local effects (ALE) plotting: `ALEPython` [Kuhn-Régnier et al., 2021a].
- Chapters 4 and 5
  - `python-inferno` [Kuhn-Régnier, 2022b].
  - `jules_output_analysis` [Kuhn-Régnier, 2022a].
  - Customised Differential Evolution Adaptive Metropolis (DREAM) implementation: `SPOTPY` [Kuhn-Régnier et al., 2022b], which was based on `SPOTPY` [Houska et al., 2015, 2022].
  - Sobol sensitivity analysis (SA): `SALib` [Kuhn-Régnier et al., 2022a], which was based on the existing `SALib` implementation [Herman and Usher, 2017, Usher et al., 2021, Iwanaga et al., 2022].

The above code made use of the `Python` [Van Rossum and Drake, 2009] packages `SciPy` [Virtanen et al., 2020], `Matplotlib` [Hunter, 2007], `NumPy` [Oliphant, 2006], `Iris` [Met Office, 2010], `Dask` [Dask Development Team, 2016], and `Jupyter` notebooks [Kluyver et al., 2016]. Fourth-generation global fire emissions database (GFED4) data were read using `pyhdf`, which wraps `NCSA HDF` version 4.

Contents

1 Introduction **19**

1.1 Significance of wildfires . . . 20

1.2 Wildfire history . . . 22

1.3 Wildfire drivers . . . 26

1.3.1 Fuel . . . 28

1.3.2 Fire regimes . . . 29

1.4 Data . . . 32

1.4.1 Remotely sensed BA datasets . . . 33

1.4.2 Fuel observations . . . 36

1.5 Fire–vegetation models . . . 37

1.5.1 The INFERNO model . . . 40

1.5.1.1 Calculation of BA in INFERNO . . . 41

1.5.2 Model evaluation . . . 44

1.6 Thesis structure . . . 47

2 Analysis Methods **50**

2.1 Emergent vs underlying relationships . . . 50

2.2 Generalised additive models . . . 51

2.3 Random Forests . . . 51

2.3.1 N-fold cross-validation . . . 53

2.4 Machine learning model validation . . . 53

2.4.1 Spatial autocorrelation . . . 54
2.5 Visualising underlying relationships using ALE plots . . . 55
2.5.1 Visualising inhomogeneities . . . 56

3 Empirical Analysis of Burnt Area Drivers 57

3.1 Introduction . . . 59
3.2 Methods . . . 60
3.2.1 Data . . . 60
3.2.1.1 Antecedent predictor variables . . . 64
3.2.2 Data processing . . . 65
3.2.2.1 Gap filling . . . 65
3.2.2.2 Interpolation . . . 67
3.2.3 Machine learning experiments . . . 67
3.2.4 Measuring predictor importance and relationships . . . 73
3.3 Results . . . 75
3.3.1 Model performance . . . 75
3.3.2 Importance of predictors . . . 79
3.3.3 Models with fewer predictors . . . 80
3.3.3.1 Importance of antecedent fuel-related predictors . . . 84
3.3.4 Current and antecedent relationships with BA . . . 85
3.3.4.1 Consistency of relationships . . . 87
3.3.4.2 Interactions . . . 90
3.3.4.3 Geographically varying timescales of importance . . . 91
3.4 Discussion . . . 94
3.4.1 Model performance . . . 94
3.4.2 Importance of predictors . . . 95
3.4.3 Underlying relationships with BA . . . 99
3.5 Conclusions . . . 101

4 Development of a New INFERNO Flammability Parametrisation 104

4.1 Introduction . . . 106

4.1.1 Application of previous results to a new flammability parametrisation . . 106

4.2 Methods . . . 107

4.2.1 Optimised metrics . . . 107

4.2.1.1 Previous model implementation sources . . . 109

4.2.2 Available JULES data . . . 111

4.2.3 Dryness parametrisation . . . 111

4.2.3.1 Dry-day period . . . 111

4.2.3.2 VPD and precipitation . . . 112

4.2.4 Fuel build-up parametrisation . . . 114

4.2.4.1 Antecedent NPP . . . 115

4.2.4.2 Litterfall soil carbon pool model . . . 115

4.2.5 Model structure . . . 116

4.2.6 Cropland parametrisation . . . 120

4.2.7 Parameter optimisation . . . 121

4.2.8 Model analysis . . . 122

4.2.8.1 Null-model analysis . . . 122

4.2.8.2 Regional analysis . . . 123

4.3 Results . . . 125

4.4 Discussion . . . 128

4.5 Conclusions . . . 133

5 Sensitivity Analysis 135

5.1 Introduction . . . 137

5.1.1 Verification of the underlying relationships using GAMs . . . 138

5.1.2 Global sensitivity analysis using Sobol SA . . . 138

5.1.3 Reduced complexity SINFERNO models . . . 139

5.2 Methods . . . 140
5.2.1 GAM analysis . . . 140
5.2.2 Sobol sensitivity analysis . . . 141
5.2.2.1 MCMC simulation . . . 142
5.2.2.2 Sobol SA grouping . . . 143
5.2.3 Reduced complexity models . . . 145
5.3 Results . . . 148
5.3.1 GAM analysis . . . 148
5.3.2 Sobol sensitivity analysis . . . 150
5.3.3 Reduced complexity models . . . 160
5.4 Discussion . . . 163
5.4.1 GAM analysis . . . 163
5.4.2 Sobol sensitivity analysis . . . 164
5.4.3 Reduced complexity models . . . 167
5.5 Conclusions . . . 169

6 General Discussion and Conclusions **172**

6.1 Summary . . . 172
6.2 Future work . . . 177
6.3 Wider implications . . . 180

List of Figures

1.1 Overview of major fire impacts on the Earth system. 20

1.2 Average combined controls on fire activity. 26

1.3 Simulated and predicted contemporary burnt fraction. 31

1.4 Processes leading from fire starts to BA. 38

1.5 Underlying processes included in the INFERNO fire model. 40

1.6 The fire-related linkages between different components of the UKESM. 42

1.7 Sensitivities of BA to herbaceous vegetation cover for different models. 44

1.8 Seasonality of simulated and observed BA. 45

1.9 IAV of simulated and observed BA for 2001–2012. 46

2.1 Random forest algorithm overview. 52

3.1 Map plot of the fraction of filled samples for FAPAR. 66

3.2 Variogram of mean GFED4 BA. 68

3.3 Significant peak detection algorithm. 74

3.4 Global climatological R^2 scores for the RF modelling experiments. 76

3.5 Observed GFED4 BA, out-of-sample predictions of the ALL model, and relative prediction error. 77

3.6 Out-of-sample BA observations and corresponding ALL model predictions. . . . 78

3.7 Mean difference between the out-of-sample BA observations and corresponding ALL model predictions. 78

3.8 PFI importances for the 15VEG_FAPAR model. 79

3.9 B-LOO CV for the 15VEG_FAPAR model. 80
3.10 Transformed variable importance metrics (Gini, PFI, SHAP, LOCO) for the ALL, ALL_NN, and 15VEG_FAPAR models. 81
3.11 Pearson correlations between all variables used in the analysis for the time period from January 2010 to April 2015. 82
3.12 Mean change in out-of-sample prediction error between the CURR and ALL models relative to mean observations. 84
3.13 First-order ALEs for different antecedent (<1 yr) relationships with FAPAR and DD in the 15VEG_FAPAR model. 86
3.14 Timescales of influence of FAPAR and DD on BA. 87
3.15 Spatial distribution of individual peak combinations for SHAP values. 88
3.16 First-order ALEs for different lags (<1 yr) from all relevant modelling experiments for the relationships between BA and FAPAR and BA and LAI. 89
3.17 First-order LAI ALEs for different lags (<1 yr) from all relevant modelling experiments for the relationships between BA and SIF and BA and VOD. 90
3.18 First-order ALE plots showing the effect of FAPAR and the 3-month antecedent dry-day period on BA in the 15VEG_FAPAR and 15VEG_FAPAR_MON models. 91
3.19 Sorted variable importance metrics (Gini, SHAP, PFI, LOCO) for the ALL model. 92
3.20 Second-order ALE plot showing the combined zeroth order (mean), first order, and second order modelled effects of FAPAR and FAPAR 1M on BA from the 15VEG_FAPAR model. 93
3.21 Second-order ALE plot showing the combined zeroth order (mean), first order, and second order modelled effects of DD and FAPAR on BA from the 15VEG_FAPAR model. 93

4.1 Transformations applied to BA data. 108
4.2 Existing INFERNO relationships. 110
4.3 Dry-day period calculation threshold estimation. 111

4.4 Comparison of JULES FAPAR and NPP to observed FAPAR. 114
4.5 Generalised sigmoid curves. 117
4.6 Mega PNV biomes. 123
4.7 Modelled and observed BA. 124
4.8 Phase differences between modelled and observed BA. 126
4.9 Average NME errors of the different SINFERNO parametrisations. 127
4.10 Average MPD errors of the different SINFERNO parametrisations. 127
4.11 Regional MPD errors of the different SINFERNO parametrisations. 129
4.12 Regional NME errors of the different SINFERNO parametrisations. 130

5.1 MCMC convergence. 142
5.2 Iterative optimisation graph. 144
5.3 Small example model configuration. 145
5.4 Most complicated model configuration. 146
5.5 Geographical distribution of stratified five-fold CV splits. 147
5.6 GAM verification of underlying relationships. 149
5.7 Sobol arcsinh-NME importances for SINFERNO-B2 data. 154
5.8 Sobol MPD importances for SINFERNO-B2 data. 155
5.9 Sobol arcsinh-NME importances for SINFERNO-B2 temperature and NPP parameters. 156
5.10 Sobol MPD importances for SINFERNO-B2 temperature and NPP parameters. . 157
5.11 Sobol arcsinh-NME importances for SINFERNO-B2 dryness and fuel parameters. 158
5.12 Sobol MPD importances for SINFERNO-B2 dryness and fuel parameters. 159
5.13 Iterative model optimisation performance. 161

List of Tables

3.1 Characteristics of the datasets used. 62

3.2 RF modelling experiments. 71

3.3 Variables used in the RF modelling experiments. 72

3.4 Ranked importance of variables in the RF experiments. 83

4.1 The different SINFERNO models. 107

4.2 Possible discrete parameter values for the 'B' dryness parametrisation. 113

4.3 Possible discrete parameter values for the '1' fuel build-up parametrisation. . . . 115

4.4 Possible discrete parameter values for the '2' fuel build-up parametrisation. . . . 116

4.5 Number of parameters being optimised across the SINFERNO models. 122

5.1 Global arcsinh-NME Sobol SA for SINFERNO-B2 input data variables. 150

5.2 Global MPD Sobol SA for SINFERNO-B2 input data variables. 151

5.3 Global arcsinh-NME Sobol SA for SINFERNO-B2 parameters. 151

5.4 Global MPD Sobol SA for SINFERNO-B2 parameters. 152

5.5 Proportion of mean global parameter importance by type measured using Sobol S_T indices for SINFERNO-B2. 153

5.6 Best-performing number of optimised parameters. 162

6.1 BEST15 RF model variables from Chapter 3. 173

6.2 Top two data variables for the SINFERNO-B2 fire model. 174

6.3 Top 10 SINFERNO-B2 fire model parameters. 175

List of Acronyms

ABC approximate Bayesian computation

AF active fire

AGB above-ground biomass

AIC Akaike information criterion

ALE accumulated local effects

AVHRR Advanced Very High Resolution Radiometer

BA burnt area

B-LOO CV buffered leave-one-out cross-validation

CCI Climate Change Initiative

CG cloud–ground

CV cross-validation

DAFI database of anthropogenic fire impacts

DD dry-day period

DGVM dynamic global vegetation model

DPM decomposable plant material

DREAM Differential Evolution Adaptive Metropolis

DTR diurnal temperature range

ECMWF European Centre for Medium-Range Weather Forecasts

ENSO El Niño/Southern Oscillation

ERA5 ECMWF Reanalysis 5th Generation

ESA European Space Agency

ESM Earth system model

FAPAR fraction of absorbed photosynthetically active radiation

FireMIP fire model intercomparison project

FRI fire return interval

FRP fire radiative power

GAM generalised additive model

GFED4s fourth-generation global fire emissions database with small fire boost

GFED4 fourth-generation global fire emissions database

GFED global fire emissions database

GHG greenhouse gas

GLM generalised linear model

Glob-FIRM GLOBal FIRe Model

GPP gross primary productivity

GPU graphics processing unit

HDI Human Development Index

HYDE History database of the Global Environment

IAV inter-annual variability

ICE individual conditional expectation

INFERNO INteractive Fire and Emission algoRithm for Natural envirOnments

JULES Joint UK Land Environment Simulator

LAI leaf area index

LIS Lightning Imaging Sensor

LOCO leave-one-column-out

LSM land surface model

MaxT maximum temperature

MCMC Markov chain Monte Carlo

MERIS Medium Resolution Imaging Spectrometer

ML machine learning

MODIS MODerate-resolution Imaging Spectroradiometer

MOSES Met Office Surface Exchange Scheme

MPD mean phase difference

MSE mean squared error

NASA National Aeronautics and Space Administration

NDVI normalised difference vegetation index

NLDI Night Light Development Index

NME normalised mean error

NMSE normalised mean square error

NPP net primary productivity

OOB out-of-bag

PDP partial dependence plot

PFI permutation feature importance

PFT plant functional type

PNV potential natural vegetation

POPD population density

REML restricted maximum likelihood

RFECV recursive feature elimination with cross-validation

RF random forest

RH relative humidity

RPD Reading Palaeofire Database

RPM resistant plant material

SA sensitivity analysis

SHAP SHapley Additive exPlanations

SIF solar-induced fluorescence

SIMFIRE simple global fire model

SINFERNO Sigmoidal INFERNO

SOFIA Satellite Observations to predict FIre Activity

SSE sum of squares error

SWI soil water index

UKCA United Kingdom Chemistry and Aerosols

UKESM UK Earth System Model

VIIRS Visible Infrared Imaging Radiometer Suite

VODCA VOD Climate Archive

VOD vegetation optical depth

VPD vapour pressure deficit

WGLC WWLLN Global Lightning Climatology

WUI wildland-urban interface

WWLLN World Wide Lightning Location Network

Chapter 1

Introduction

Wildfires are a significant natural disturbance with profound impacts on the Earth system, affecting vegetation, atmospheric composition, and climate [Bowman et al., 2011, Voulgarakis and Field, 2015, Andela et al., 2017, Lasslop et al., 2019]. Wildfires also affect humans both indirectly and directly through their significant socioeconomic impacts in the form of loss of life and property damage [Duff and Penman, 2021] and deterioration of health due to air quality degradation [Johnston et al., 2012, Marlier et al., 2013, Johnston and Bowman, 2014, Johnston, 2017, Roberts and Wooster, 2021], amongst others. How fire regimes may change in the future and how fire-related feedbacks may therefore influence climate and environmental changes are growing concerns. Improving our understanding of the factors that influence wildfires and thereby enhancing our ability to predict fire regime shifts due to global warming, is therefore a key goal of this thesis.

This thesis first discusses the significance of wildfires and what drives them on a global scale. Machine learning techniques are then used to explore the controls (drivers) of wildfires in an empirical modelling framework. Thereafter, the insights gained from these analyses are used to carry out an investigation into how fire drivers can be modelled computationally and how the INteractive Fire and Emission algoRithm for Natural envirOnments (INFERNO) fire model can be modified in order to improve wildfire predictions. Finally, multiple global sensitivity analysis approaches are used to investigate the sensitivity of the improved fire model parametrisation to

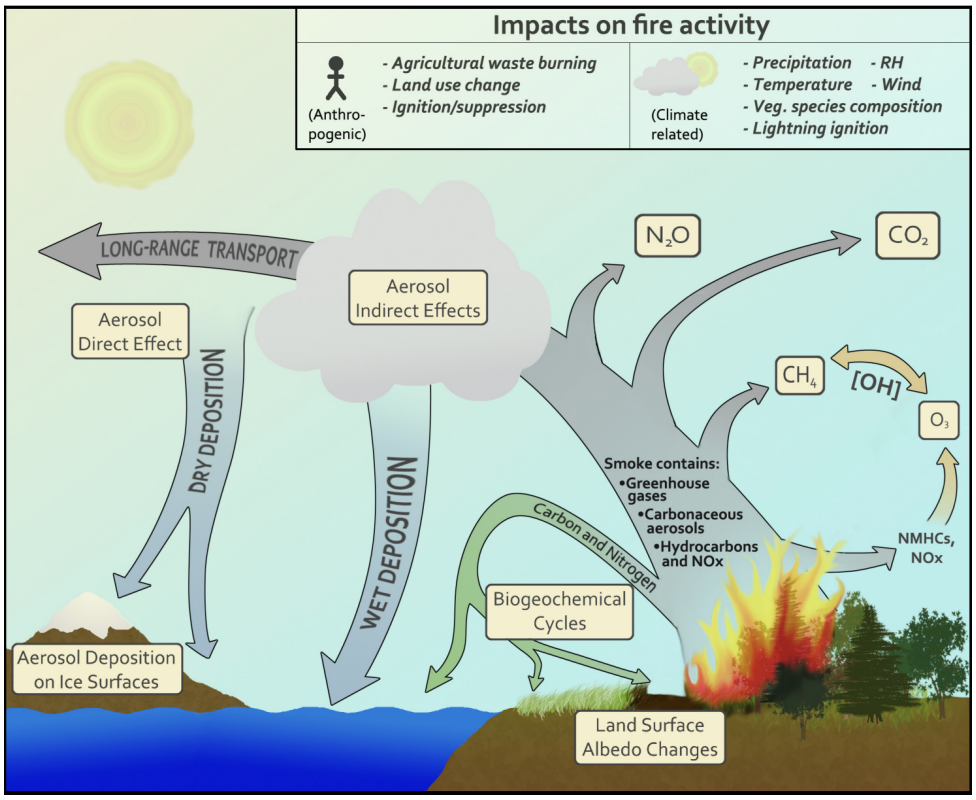


Figure 1.1: Overview of major fire impacts on the Earth system. Source: Fig. 1 from Ward et al. [2012], licensed under CC BY 3.0.

different input data and parameters.

In this chapter, I will introduce the importance of wildfires along with their historical significance. Additionally, the drivers of fire regimes (including human interactions), the data needed to investigate wildfires, and how they can be modelled will be discussed. Finally, I will present the structure of the remainder of the thesis.

1.1 Significance of wildfires

Wildfires are a naturally occurring process with global ecological and socioeconomic implications (see Fig. 1.1). As will be discussed further below, they are responsible for a significant fraction of inter-annual variability (IAV) of CO_2 and aerosol emissions [Keppel-Aleks et al., 2014, Voulgarakis and Field, 2015, Voulgarakis et al., 2015] and are involved in shaping and maintaining the distribution of vegetation types. It has been claimed that humans have used wildfires as a tool for over 40 000 years [Bowman et al., 2011, Scott, 2018], while the domestication of fire has been

called "a defining trait of humanity" by Bowman et al. [2011].

Wildfires emit the equivalent of around 25 % of fossil fuel CO_2 emissions [Hantson et al., 2016]. However, post-fire vegetation regrowth can absorb the carbon initially emitted during wildfires. As such, these emissions are usually reabsorbed within a "disturbance-recovery episode" [Yin et al., 2020]. In such a dynamic equilibrium state, wildfires thereby contribute more to the IAV of atmospheric CO_2 than to its overall concentration. The degree of reabsorption and its timescale depend on the vegetation present, with savannahs and grasslands typically reabsorbing the emitted carbon within the next growing season [Sommers et al., 2014], for example. However, in the presence of land-use changes (e.g. from forested areas to grasslands), post-fire regrowth may no longer be able to completely absorb the emitted carbon as fire regimes are affected. For example, man-made fire use is currently implicated in the large-scale conversion of carbon-rich forested areas to pastures or croplands that have a far lower potential to store carbon, thereby altering the way in which emissions are reabsorbed during post-fire vegetation regrowth [Andela et al., 2017]. Such land-use changes, which alter the terrestrial carbon sink, will then result in a net increase in atmospheric greenhouse gas (GHG) levels [Andela et al., 2017].

Furthermore, wildfires alter the atmospheric composition by emitting other gases, e.g. carbon monoxide (CO), nitrogen oxides (NO_x), and non-methane volatile organic compounds (NMVOCs), as well as aerosols like black carbon (BC) and organic carbon (OC) [Voulgarakis and Field, 2015, van der Werf et al., 2017]. Air pollution and climate are affected by emitted carbonaceous compounds (BC, OC, CO, NMVOCs, and methane (CH_4)), and NO_x, NMVOCs, CO, and methane also act as ozone precursors in the troposphere, thereby contributing to ozone levels without their direct emission by fires themselves [Voulgarakis and Field, 2015].

The ecological impacts of fires include the maintenance of specific vegetation types. For example, the dominance of fire-adapted C4 grasses in fire-prone savannah biomes has been attributed to their high productivity in post-fire environments compared to C3 grasses [Robinson et al., 1979, Bond et al., 2003, Tix and Charvat, 2005, Ratnam et al., 2011, Ripley et al., 2015]. This relatively easier regrowth allows C4 grasses to outcompete C3 grass species following a fire

disturbance [Ripley et al., 2015]. The extent of this fire maintenance is such that it has been estimated that global forest cover would roughly double in the absence of fire [Bond et al., 2005]. However, this is a model-dependent estimate, and more recent estimates using an ensemble of four global vegetation–fire models from the fire model intercomparison project (FireMIP) have shown that the reduction of tree cover due to fire may be up to around 40 % [Rabin et al., 2017, Lasslop et al., 2020]. Wu et al. [2022] also carried out future simulations using the fire-enabled LPJ-SEVER model—based on the LPJ-DGVM—to show that the decreasing fire trends that are predicted under future demographic and climatic scenarios are expected to lead to increased global woody encroachment (increase in globally averaged tree cover) by almost 5 % compared to a constant-fire scenario. Specific vegetation traits have also emerged as a response to wildfires in other biomes, such as resprouting or the thicker bark of certain trees [e.g. Harrison et al., 2021], a process that has been estimated to have occurred up to 100 million years ago during the Cretaceous period [Scott, 2018].

1.2 Wildfire history

Fires have been present ever since plants established themselves on land [Scott, 2018], and plant species diversification has gone hand in hand with the diversification of fire regimes in the past. Future fire activity and the impact of fire-related feedbacks on the global climate and environment are increasingly relevant owing to humanity's increasing impact but are best viewed within its historical context. As mentioned previously, humans have long used fire for a variety of purposes ranging from hunting, fire regime management, and farming to cooking and metalwork, amongst others. However, the human relationship with fire has seen dramatic changes over the last thousands of years [Bowman et al., 2011]. According to the database of anthropogenic fire impacts (DAFI) [Perkins et al., 2021], current anthropogenic fire usage is global and varied, with uses ranging from prescribed burning (pyrome management) to agriculture (e.g. crop residue burning) and arson.

Using evidence of fire over the past millennia, Marlon et al. [2008] and Wang et al. [2010]

showed that fire activity has declined over the last century, likely caused by increased agricultural land use, leading to reduced fuel loads and reduced fuel connectivity [Prentice, 2010]. This decreasing trend in global burnt area (BA) may also be apparent throughout the last decades [van Lierop et al., 2015, Doerr and Santín, 2016, Andela et al., 2017], based on remotely sensed BA data like the global fire emissions database (GFED) [Giglio et al., 2013], although Forkel et al. [2019b] showed that the short observational time period that is currently available from such remote sensing data significantly limits our ability to make robust claims about recent global BA trends.

The aforementioned recent trend in global BA is far from being globally homogeneous, however, with distinct patterns emerging depending on the region. For example, there is an increasing trend in North American boreal BA, with more large fires and an increased proportion of lightning ignitions compared to human ignitions [Kasischke and Turetsky, 2006]. Climate change is likely responsible for drier conditions and resulting increases in the fire season length, which consequently led to increased BA [Kasischke and Turetsky, 2006]. This is supported by other studies, which found a potential positive feedback on further boreal BA and carbon emissions as a result of the climate-change-induced increases in boreal BA [Oris et al., 2014, Zheng et al., 2021].

Recent declines in BA in some regions have been explained as a consequence of human activity through indirect and direct intervention [Martínez et al., 2009, Andela et al., 2017], albeit modulated by climate and vegetation type [Forkel et al., 2019b]. In South America, empirical analysis has demonstrated that the importance of anthropogenic factors, mostly the influence of land conversion, outweighs climatic fire drivers [Aldersley et al., 2011]. Andela et al. [2017] claimed that increasing agricultural activity, especially in African savannahs, has contributed greatly to a decrease in global BA by $24.3 \pm 8.8\,\%$ between 1998 and 2015. With $64\,\%$ of global BA occurring in Africa, it is of major global importance [Grégoire et al., 2012]. Land use conversion to cropland in Africa has the effect of reducing BA because agricultural areas have lower fuel levels for most of the growing season (sometimes no fuel at all), highly controlled fire use (often

used sparingly to burn stubble and crop residue), and strong fire suppression [Grégoire et al., 2012]. Compared to cropland, fires in natural vegetation also tend to be larger, as e.g. large fires are expected in the open savannah regions of Africa [Grégoire et al., 2012].

The observed expansion of cropland and pasture conversion in the south American Amazonia and Cerrado regions has been increasingly scrutinised, with fire-driven land-use conversion starting in the late 1980s [see e.g. Davidson et al., 2012, van Marle et al., 2017a, Noojipady et al., 2017, Zalles et al., 2019, Hong et al., 2021]. In addition to dry season wildfires, fire is used for land clearing and pasture management [Artaxo et al., 2002, Davidson et al., 2012]. One of the major drivers of the observed land-use change is the planting of soybean crops [Morton et al., 2006, Gibbs et al., 2015, Fehlenberg et al., 2017, Zalles et al., 2019], with some studies finding a correlation between forest clearance for intensive mechanised agriculture and soybean prices [Morton et al., 2006]. Carbon loss from areas deforested for croplands may also be greater than for pastures due to the rapid and complete removal of above-ground biomass (AGB) and woody roots (using fire) and little to no carbon offset from subsequent crop planting [Morton et al., 2006]. Due to increased anthropogenic fire suppression as a result of an increase in tropical cropland area by almost 80 % and an increase in the tropical population by over 300 % from the 1900s to the 2000s, human activities likely led to an overall decrease in tropical BA despite increased deforestation rates in Amazonia since the 1970s [Houghton et al., 2000, Yang et al., 2014]. And while an analysis of Amazonian deforestation dynamics between 2001 and 2014 revealed reduced deforestation rates overall, it was also found that the proportion of small deforestation events increased [Kalamandeen et al., 2018].

However, Forkel et al. [2019b] demonstrated, using the fourth-generation global fire emissions database (GFED4) and FireCCI50 satellite BA products, that the short observational time period currently available for such satellite data significantly hampers our ability to make robust claims about global burnt area trends from such data, since the sign and significance of the global BA trend depend on the start and end dates chosen for the analysis. Also, while Andela et al. [2017] reported similar trends for both the fourth-generation global fire emissions database with small

fire boost (GFED4s) [Randerson et al., 2012, van der Werf et al., 2017] and MCD64A1 [Giglio et al., 2018] datasets, such datasets may be subject to significant uncertainties (see e.g. Zhang et al. [2018] for GFED4s), as discussed further in Section 1.4.1.

Human intervention can promote or suppress fire through ignitions, fuel management, and landscape modification [Bowman et al., 2011]. Furthermore, humans indirectly influence fire through their interference with global and regional climate. Direct human influence was found to be the major driver of changes in fire regimes [Morgan et al., 2001] in recent decades [Martínez et al., 2009, Andela et al., 2017]. Land use changes are also implicated in ignitions and suppression dynamics (e.g. via agricultural abandonment and development [Martínez et al., 2009] or landscape fragmentation [Archibald et al., 2009]). But the predominance of human influence on fire regimes may change in the future as climate change affects fire regimes, and consequently, climate is expected to play an ever-increasing role once again [Pechony and Shindell, 2010]. Using a process-based fire model, Yang et al. [2014] corroborated these findings by predicting that, given continued warming in the 21st century, climate may become more significant than anthropogenic impacts with respect to wildfires in the extratropics [Yang et al., 2014]. A mainly temperature-driven increase in conditions conducive to wildfires was also suggested by a number of regional studies [e.g. Westerling, 2006, Barbero et al., 2015, van Oldenborgh et al., 2020, Goss et al., 2020], with temperature extremes akin to those causing the massive 2019/20 Australian wildfires having become more likely due to global warming by at least a factor of two [van Oldenborgh et al., 2020]. On a global scale, Abatzoglou et al. [2019] showed that anthropogenic climate change had led to an increase in fire weather over 22 % of the global burnable area by 2019, while Jolly et al. [2015] found that anthropogenic climate change has led to a lengthening of the fire season across more than a quarter of global vegetated land in recent decades. Increases in fire weather are also predicted under different assumptions about levels of future warming [e.g. Bedia et al., 2015, Bowman et al., 2017, Burton et al., 2018, Turco et al., 2018].

With extreme climate events seemingly becoming ever more common [Herring et al., 2015], extreme wildfire events (e.g. in Australia [Sanderson and Fisher, 2020, van Oldenborgh et al.,

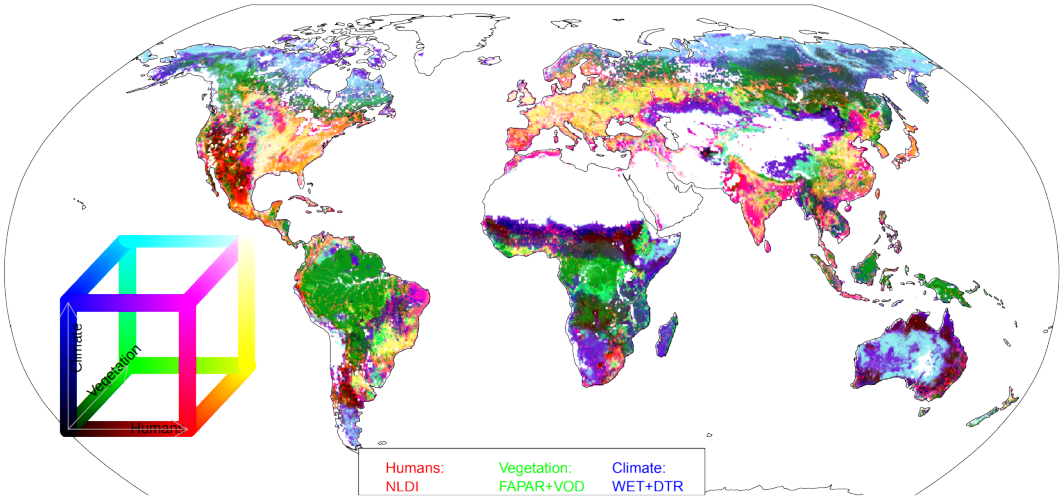


Figure 1.2: Fire activity restrictions coloured according to the prevalence of three categories. Human influence is shown in red and measured using the Night Light Development Index (NLDI), a proxy for human development. Vegetation is shown in green and measured using a combination of one-month precedent FAPAR and annual VOD, two vegetation proxies. Climatic influence, shown in blue, is measured using the monthly number of wet days (WET) and the diurnal temperature range (DTR). Bright colours indicate a restriction of burning. Colour combinations indicate restriction combinations, with e.g. yellow indicating limitation by both humans and vegetation. Source: Fig. 6f from Forkel et al. [2017], licensed under CC BY 3.0.

2020], Canada [Kirchmeier-Young et al., 2019], and the US [Yoon et al., 2015]) have recently been increasingly scrutinised from a climate change perspective [Yoon et al., 2015]. Indeed, wildfires in some areas of the world have increased, possibly as a result of global warming, thereby increasing the damage caused as well as firefighting efforts and expenditures [Westerling, 2006].

1.3 Wildfire drivers

While most of the focus for future fires has been on climate factors, it is well-known that other factors are important drivers too. Generally, the factors that influence the occurrence and intensity of fires are the presence of an ignition source, vegetation properties that determine the fuel quantity and quality, and weather conditions that promote fuel drying and fire spread [Bowman et al., 2009, Krawchuk et al., 2009, Moritz et al., 2012, Knorr et al., 2014, Bistinas et al., 2014, Forkel et al., 2017, 2019b, Jones et al., 2022]. Strong surface winds, dry fuel, and the presence of an ignition source such as lightning create ideal wildfire conditions [Krawchuk and Moritz, 2011]. Human activities are also implicated, either promoting or suppressing fire through

ignitions, fuel management, and landscape modification [Bowman et al., 2011]. In all but the most remote regions, natural lightning ignitions are usually insignificant [Jones et al., 2022]. While several regions, e.g. boreal regions or southeast Australia, do experience a significant proportion of lightning ignitions, most do not, as e.g. tropical lightning typically coincides with wet conditions [Jones et al., 2022]. The factors driving wildfires are strongly coupled to one another: climate conditions influence the incidence of lightning and also the type of vegetation, while wind strength and the impact of atmospheric conditions on drying are modulated by vegetation cover [Harrison et al., 2021]. It is currently unclear how the increasing temperatures due to climate change will impact lightning occurrence since the expected decrease in ice formation will inhibit lightning, while the expected increase in convection rates would be expected to increase it [Jones et al., 2022]. As such, the expected influence of climate change on lightning occurrence currently remains model-dependent [Jones et al., 2022]. A representation of the resulting complex variation of different contemporary controls on wildfire activity is shown in Fig. 1.2.

The relationships between ignitions, vegetation, and climate may depend on the timescales involved: short-term drought promotes fuel drying and hence increases the risk of fire, but longer-term drought conditions will reduce vegetation growth and fuel loads. These drivers may also have different, at times antagonistic, effects on different aspects of the fire regime [Archibald et al., 2009]. For example, human ignitions can, amongst others, affect the number of fires (e.g. for southern California [Keeley et al., 1999]) and their seasonal distribution [Le Page et al., 2010, Benali et al., 2017], while having little impact on fire spread and the total area burnt [Syphard et al., 2007]. This is because depending on the region (e.g. in South Africa [Archibald et al., 2009]), increased human settlement can increase land management, landscape fragmentation, and fire suppression (especially at the wildland-urban interface (WUI) [Syphard et al., 2007]), which then impedes fire spread, cancelling out the potential for increased BA due to the increased number of ignitions. How frequently wildfires occur and how much area is burnt depends on many factors, most notably the type of vegetation present, climate conditions such as the fuel moisture and wind, the available biomass, anthropogenic interactions with fires, and natural ignition

sources such as lightning [Krawchuk and Moritz, 2011, Moritz et al., 2012]. The differences between controls on different aspects of fire regimes have also been demonstrated in a global study by Haas et al. [2022]. They used generalised linear models (GLMs) to analyse BA, fire spread, and fire radiative power (FRP) together with a range of possible wildfire drivers to find that BA is predominantly driven by fuel availability and dryness and ignitions while being constrained by landscape fragmentation. As expected, they found fire size to be driven by conditions that promote fire spread, such as wind speed and dryness (as measured by the diurnal temperature range (DTR) and vapour pressure deficit (VPD)), while also being limited by landscape fragmentation. This limiting influence of VPD is expected, as it has been found to limit plant growth [Fu et al., 2022]. Fire intensity was found to be driven by fractional tree cover and road density and constrained by fuel moisture (as measured by VPD and dry-day period (DD)) and ignitions.

1.3.1 Fuel

The tendency to burn is significantly affected by fuel characteristics such as fuel type, amount, and arrangement [Archibald et al., 2009]. Fuel can be defined as organic matter that is available for ignition and later combustion [Keane et al., 2001]. Its characteristics can be highly variable, with fuel types ranging from aerial crown fuels to surface fuels, with both dead and live fuel potentially playing a role in combustion [Keane et al., 2001]. Since the litter layer is where most fires are initiated, its wetting and drying are particularly important [Venevsky et al., 2002]. Smouldering fires in peat also play a large role in certain regions [Turetsky et al., 2011, Sommers et al., 2014, Huijnen et al., 2016, Atwood et al., 2016, van der Werf et al., 2017, Scholten et al., 2021].

Along with the varying types of fuel, its properties, such as shape, size, and density—influenced by vegetation, weather, and humans—in turn affect fire spread, in addition to how rapidly fuel dries out and becomes combustible. Antecedent weather conditions in the weeks to years before fire events can determine fuel availability [van Oldenborgh et al., 2020] and hence fire occurrence

and extent, as the accumulation of fuel over time will influence fire intensity. However, the effect of antecedent weather conditions on BA may depend on the types of vegetation present (which influence whether fuel drying or accumulation is most important): antecedent precipitation will increase BA in fuel-limited regions but decrease BA in regions where fuel drying is the major control [Littell et al., 2009, Abatzoglou and Kolden, 2013, Alvarado et al., 2020]. Moreover, human impacts on fuel structure are currently not well-understood [Archibald et al., 2018]. Prediction of BA, a key goal of this thesis, will require an understanding of the timescales of fuel accumulation, the interaction between biophysical drivers and fuel build-up, and the effects of antecedent weather conditions on both fuel loads and fuel drying. While it has been observed that catastrophic mega-fires are becoming more common following human fire exclusion and the resulting fuel build-up on decadal timescales or longer [e.g. in the USA; Marlon et al., 2012, Higuera et al., 2015, Parks et al., 2015], an empirical analysis will necessarily be focused on approximately yearly timescales of fuel build-up since the currently available global observational record cannot capture the decadal-or-longer fire return intervals (FRIs) observed e.g. in temperate or boreal forests.

1.3.2 Fire regimes

The type of fire determines the effects on vegetation, with changing fire sizes and frequencies (FRI) affecting contemporary vegetation structures [Andela et al., 2017]. Such effects differ both temporally and spatially. Small fires can trigger gap dynamics on a local scale while also opening canopies, whereas the expression of fire-related trait syndromes is strongly controlled by the occurrence and nature of large fires [Harrison et al., 2021]. Thus, the prevalence of wildfires is also implicated in the maintenance of specific vegetation types, notably savannahs and grasslands [Bowman et al., 2011]. Such landscapes are home to fire-adapted plants, e.g. with fire-promoted germination or a large proportion of biomass below ground and thereby shielded from the heat of frequent fires [Pausas et al., 2018, Simpson et al., 2019]. Vegetation structure is also influenced by wildfires, with some plant species evolving to display increased resilience

to fire disturbance using fire-dependent seed germination to enable post-fire regeneration, for example [Paula and Pausas, 2008, Moreira et al., 2010].

Wildfires can be categorised into fire regimes based on their frequency (FRI), size (BA), intensity (fuel consumption), and the season they occur in [Morgan et al., 2001, Moritz, 2003, Bond and Keeley, 2005, Whitlock et al., 2010, Archibald et al., 2013, Hempson et al., 2018]. Fire regimes (pyromes) are not necessarily related to their underlying vegetation (biomes), however [Archibald et al., 2013]. Therefore, changes in vegetation may not be directly linked to changes in fire regimes. At least in part, this perhaps counter-intuitive decoupling may be due to human influences that both act to delineate pyromes while also forming an entirely separate global pyrome that is present across vegetation types and climate zones [Archibald et al., 2013]. This separate pyrome is marked by small fires, likely as a result of physical barriers to fire spread that are created by man-made landscape changes [Archibald et al., 2013]. Additionally, the presence of alternative stable states may also contribute to this, i.e. the presence of different vegetation such as forest or fire-maintained grassland under the same climatic conditions primarily due to different disturbance histories [Bond et al., 2005, Archibald et al., 2009]. For example, sharp forest-savannah transitions have been explained using this approach [Staver et al., 2011, van Nes et al., 2018, Pausas and Bond, 2020]. While there is some disagreement over how widely applicable this concept is, with edaphic properties being a potential alternative explanation [e.g. Veenendaal et al., 2018], it is generally agreed that positive feedbacks between vegetation and fire exist, leading to sharper vegetation transitions as a result [Harrison et al., 2021].

These findings complicate the straightforward modelling of fire using only instantaneous biophysical properties. Harrison et al. [2021] argue that, in addition to the enduring importance of enhancing our understanding of the relationships that link vegetation to fuel build-up and drying, improved fire models should take greater advantage of our current understanding of fire ecology.

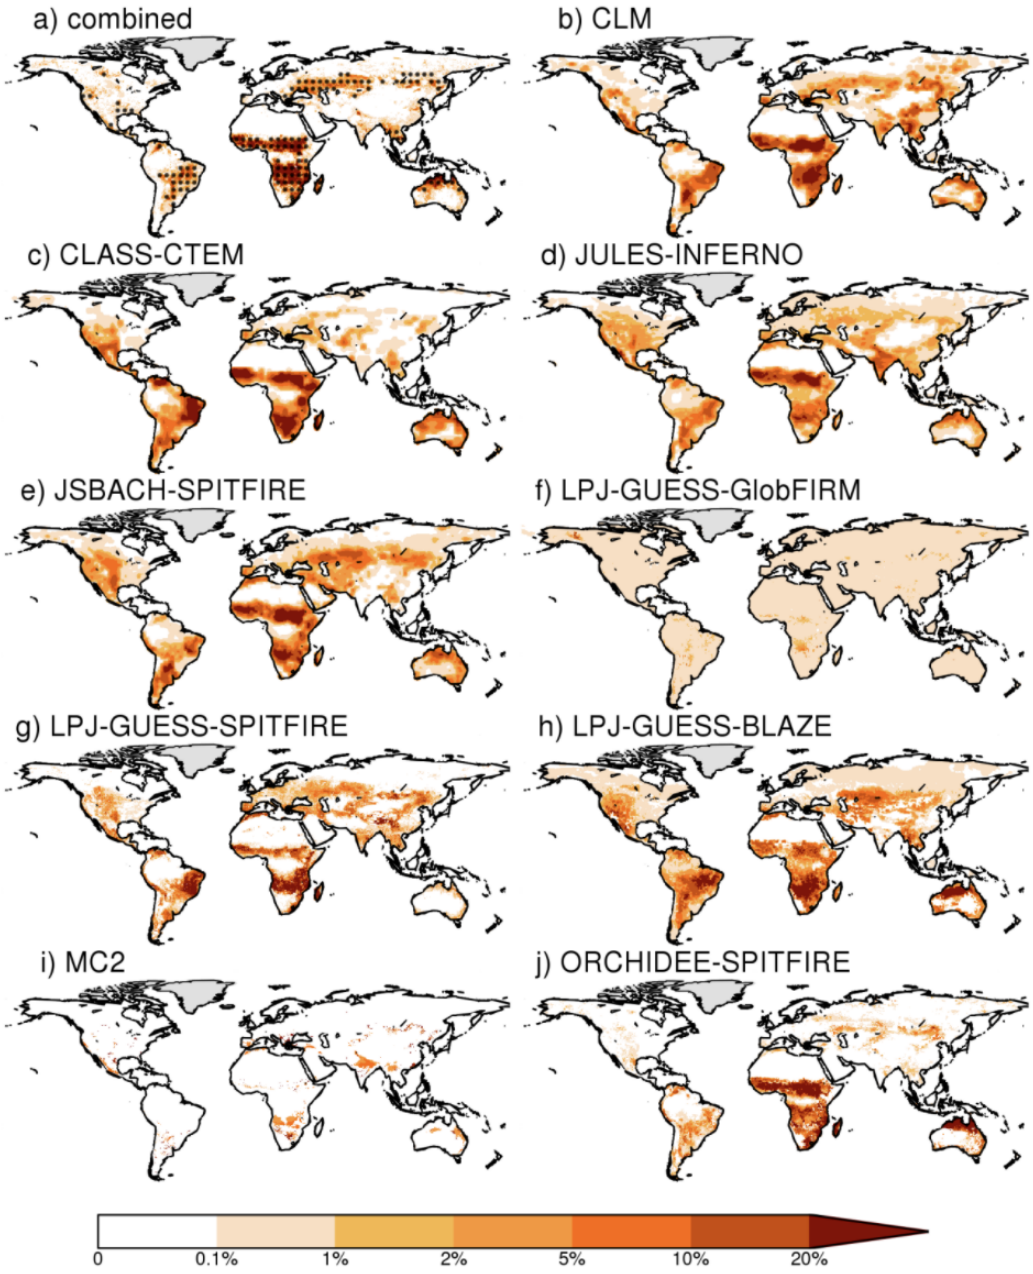


Figure 1.3: Simulated and predicted contemporary (2002–2012) burnt fraction. The combined observations (a) are the mean observations from the following BA datasets: GFED4, GFED4s, MCD45, FireCCI40, and FireCCI51. Stippled regions have a variance between all BA datasets less than the observed mean. Source: Fig. 2 from Hantson et al. [2020], licensed under CC BY 4.0.

1.4 Data

Data required for the study of wildfires include data about fire activity and its drivers, such as climate, weather, vegetation, and fuel, amongst others. Contemporary data are usually based on satellite data (remotely sensed), which are often combined with local measurements and climate models in a step known as reanalysis [see e.g. Zeng et al., 2015, Copernicus Climate Change Service (C3S), 2017, Hersbach et al., 2020], in order to yield global estimates of key properties. However, satellite observations have only been available for a few decades, meaning that other means must be used to investigate fires further into the past (e.g. using tree rings and charcoal records), with investigations of future fires relying on climate model predictions.

While palaeofire reconstructions will not be used here, they provide a useful context for more contemporary data sources and analysis of long-term carbon sinks. Past fire activity may be inferred both quantitatively and qualitatively from multiple sources (e.g. tree rings, archaeology, or written records) [Bowman et al., 2011]. However, analysis of sedimentary charcoal is the most widely used technique to reconstruct high-resolution fire activity across millennia [Iglesias et al., 2015]. Charcoal particles are formed during wildfires as a result of incomplete combustion before being preserved in sediments. Most recently, global palaeofire data based on charcoal records have been collated in the Reading Palaeofire Database (RPD), which contains 1382 records spanning the past 22 000 years with at least a resolution of 1000 years [Harrison et al., 2022]. However, despite the recent growth in the number of records, there exists a spatial bias towards the northern extratropics [Harrison et al., 2022]. Dry regions are also under-represented due to their general lack of suitable conditions for the preservation of charcoal [Leys et al., 2018, Harrison et al., 2022]. There are also multiple processes that influence the production of charcoal, its transport, and finally, its deposition. Thus, the area within which charcoal originates depends on the size of the charcoal particles; more localised analyses are possible using larger macrocharcoal particles that do not undergo widespread transport mechanisms to the same extent as smaller particles [Iglesias et al., 2015, Marlon et al., 2016, Graves et al., 2019].

1.4.1 Remotely sensed BA datasets

BA observations for the present day are predominantly based on remotely sensed reflectance (vegetation index) changes of the terrestrial surface [Chuvieco et al., 2019]. Along with active fire (AF) observations, which are derived from the thermal energy emitted by fires that burn at the time of satellite overpass, remotely sensed fire data allow the quantification of global BA (see e.g. Fig. 1.3a) and the energy released by individual fires [Wooster et al., 2021], as well as longer-term properties such as the FRI if observations are available for a long enough time period (which they generally are not, however).

While ground-based wildfire observations may be available for a longer time period than satellite observations for certain locations (e.g. local fire records between 1941–1996 in Kruger National Park in South Africa [Van Wilgen et al., 2000]), the key advantages of remotely sensed data are that all observations are made using the same instrument (thereby eliminating discrepancies that may result from location-specific monitoring techniques) and the fact that satellites may attain global coverage.

Remotely sensed global BA products have been produced for over two decades. The first such products include GBA-2000 [Tansey et al., 2004], GLOBSCAR [Simon et al., 2004], Globcarbon [Plummer et al., 2006], and L3JRC [Tansey et al., 2008]. While BA products have also been created based on the Advanced Very High Resolution Radiometer (AVHRR) sensor, which was first launched in 1978 and provides global daily coverage (albeit at a relatively coarse spatial resolution), combining such products with current MODIS-based BA products is inherently problematic due to temporal artefacts caused by e.g. orbital drift [Giglio and Roy, 2020]. The first such NASA MODIS BA product was the MCD45 product [Roy et al., 2008]. This has now been superseded by the collection 6 MCD64A1 product [Giglio et al., 2018], the previous collection 5.1 (C5.1) version of which was used as the basis for the often-used GFED4 BA product [Giglio et al., 2013]. ESA Climate Change Initiative (CCI) BA products have also been released, based both on the Medium Resolution Imaging Spectrometer (MERIS) (FireCCI31, FireCCI41) [Alonso-Canas and Chuvieco, 2015] and MODIS (FireCCI50) sensors [Chuvieco et al.,

2018]. Multiple distinct BA estimates may be calculated despite using the same underlying data (e.g. MODIS data) due to algorithmic differences, as will be illustrated further below.

Despite the advantages of global remotely sensed (satellite) wildfire datasets, they also suffer from several issues, including low temporal coverage (e.g. MODIS data are only available starting in 2000), low inter-dataset agreement, and spurious handling of small fires [Zhang et al., 2018]. The use of different algorithms and satellite sensors also leads to discrepancies between remotely sensed wildfire datasets, even if the same underlying data from the same sensor are used (see e.g. [Hantson et al., 2016, Humber et al., 2019]).

Small fires, in particular, may present a significant contribution, with an increasing proportion of small-scale deforestation events in the Amazon region [Kalamandeen et al., 2018], a 37 % increase in global BA when small fires were estimated using AF data in the GFED4s dataset [Randerson et al., 2012, van der Werf et al., 2017], and an up to 120 % increase in estimated Sub-Saharan African BA compared to MODIS BA using Sentinel-2 observations and Visible Infrared Imaging Radiometer Suite (VIIRS) AF data [Chuvieco et al., 2022]. However, while the use of AF data to capture the area burnt by small fires in the GFED4s dataset has the potential to better capture global patterns of burning, it has also been found to be affected by multiple errors, including timing inaccuracies and false positives [Zhang et al., 2018], thus introducing additional flaws to the GFED4 dataset underlying the GFED4s dataset. Since small fires are often not detected by satellites [Roteta et al., 2019, Chuvieco et al., 2022], areas with a large number of small fires could exhibit a disproportionately smaller average BA compared to regions with a relative abundance of larger fires. As global ground-based measurements are not readily available [Keane et al., 2001, Mangeon, 2017], it is also hard to judge the overall accuracy of these remotely sensed datasets.

Validation of BA products based on high-resolution images [Boschetti et al., 2009], commonly Landsat images [e.g. Tansey et al., 2008, Roy and Boschetti, 2009, Bastarrika et al., 2011, Hantson et al., 2013, Padilla et al., 2014a,b, Araújo and Ferreira, 2015, Padilla et al., 2015, Boschetti et al., 2016, Roteta et al., 2019, Humber et al., 2019, Roy et al., 2019], has the potential to remedy this, since more small fires may be observed using such products compared to e.g. coarser MODIS-

based BA products. In many regions, especially those with a high proportion of relatively small fires (i.e. <100 ha [Roteta et al., 2019], e.g. agricultural), validating MODIS- or MERIS-based BA products against Landsat imagery reveals large omission errors, since these products are not able to detect small fires (see e.g. [Hantson et al., 2013, Padilla et al., 2014b, Araújo and Ferreira, 2015]), with some studies reporting omission errors as high as 72 % [Padilla et al., 2014b].

Based on their development of high-resolution BA datasets for sub-Saharan Africa, both Roteta et al. [2019] and Chuvieco et al. [2022] predicted that global BA estimates should increase with the detail (i.e. resolution) of remotely sensed BA products since this should increase the number of small fires that can be detected (see also [Roy et al., 2019]). Both their products were able to capture at least 80 % more BA than the MODIS MCD64A1 dataset through the use of the Sentinel-2 sensor with 20 m pixels, compared to the 250–500 m pixels used by MODIS. The trade-off incurred by this higher resolution (and the resulting increased ability to detect small fires) is that, per Sentinel-2 satellite, images are only available every 10 days, while MODIS images are available daily. Note, however, that combining data from the two Sentinel-2 satellites improves the revisit interval to 5 days [Chuvieco et al., 2022]. Sufficiently small revisit intervals are especially important in tropical ecosystems, where there is frequent cloud cover, and post-fire reflectance changes are relatively short-lived [Roteta et al., 2019]. Persistent cloud cover coupled with fast regeneration of vegetation could therefore prevent the detection of the reflectance changes necessary to measure BA.

While Landsat data can capture smaller fires than MODIS, it has a 16-day repeat cycle, making cloud coverage and temporal accuracy a hindrance to the development of purely Landsat-based global BA products [Boschetti et al., 2015, Hawbaker et al., 2017, Roy et al., 2019]. However, by combining Sentinel-2 and Landsat observations, Li and Roy [2017] achieved median average revisit intervals of just 2.9 days, thereby approaching the daily resolution afforded by MODIS [Roy et al., 2019]. Notwithstanding these improvements, the challenge of validating such a BA dataset at its high spatial resolution remains since it is difficult to obtain the necessary independent data [Boschetti et al., 2009, Roy et al., 2019].

Note that in addition to the difficulty of observing small fires, Mangeon et al. [2016a] estimated that in some regions, up to 50 % of the BA in intense wildfire events may not be captured by remotely sensed datasets. However, these disagreements were likely a product of the different methods used to ascertain BA when comparing ground-based to remotely sensed observations, with unburnt patches of land in central regions of the fire, in particular, being difficult to survey from the ground without resorting to costly aircraft surveillance [Mangeon et al., 2016a].

These issues present a major challenge for empirical analyses of wildfire drivers, as the derived relationships may not be consistent when comparing small and large fires. The fact that the latter are detected much more easily and thus more frequently in global datasets (like GFED4) implies that relationships inferred from remotely sensed wildfire datasets are likely unable to represent the full spectrum of fire regimes, with the results being biased towards regions with more large fire events, where omission errors have been shown to be the lowest [Padilla et al., 2014b].

1.4.2 Fuel observations

Observation and mapping of fuel characteristics are complicated by the cost of local reconnaissance efforts and the imperfections of remotely sensed fuel datasets and modelling approaches [Keane et al., 2001]. With field data on fuel loads therefore being sparse, the only global fuel dataset [Pettinari and Chuvieco, 2016] is based on extrapolating scattered field measurements by biome. Especially concealment of the ground by the canopy [Roberts et al., 2018] or confusion of canopy-level fuel with downed fuel can negatively affect the quality of remotely sensed datasets such as fuel and landcover measurements (the latter may be used as a proxy for fuel datasets) [Keane et al., 2001]. These problems also affect other datasets reliant on observing the ground, such as soil moisture or AF retrievals [Roberts et al., 2018, Wooster et al., 2021], which are often of poor quality in heavily forested regions. While remotely sensed fuel datasets are unlikely to be helpful for local, small-scale firefighting efforts [Keane et al., 2001], they may, however, be more appropriate for large-scale studies such as the one considered herein [Archibald et al., 2009]. In

addition to spatial variability as a result of inhomogeneous vegetation, weather, and climate conditions, fuel loads can also vary on short timescales. Fuel loading (litter) on the floor can be altered by up to a factor of three due to strong wind or snow events [Keane et al., 2001].

1.5 Fire–vegetation models

There are numerous fire models with different complexities, usually falling into one of three categories: process-based, empirical (statistical modelling), or a mixture of the two, where statistical parametrisations are used for some but not all processes [Perry, 1998, Hantson et al., 2016]. The key difference between these modelling approaches is their representation of the underlying processes that drive fire activity: while process-based models explicitly account for these, empirical models approximate their resulting outcomes using averaged responses to relevant variables, e.g. temperature and precipitation for biophysical processes [Hantson et al., 2016]. These approaches can also be referred to as "top-down" (empirical) and "bottom-up" (process-based) [Hantson et al., 2016].

For example, aggregate anthropogenic influences on fire activity, one of the most uncertain aspects of fire modelling [Forkel et al., 2019a], are currently represented using empirical parametrisations in most fire models. Such parametrisations rely on relationships between the number of ignitions and the population density (POPD) or Human Development Index (HDI) in a given grid cell, for example [Teixeira et al., 2021]. This is contrary to more explicit representations of human impacts like agent-based models wherein the interactions between agents (e.g. farmers) and fire are influenced by their interactions with other agents, laws, the economy, and how they perceive the state of the environment that is simulated in their vicinity [Ford et al., 2021]. A simplified example of such a scheme is found in Pfeiffer et al. [2013], where the relationship between anthropogenic ignitions and POPD depends on the prevailing human economic system, which can be one of the following three: hunter-gatherer, pastoralists, and farmers.

Fire models have been used traditionally for a long time to aid local fire management efforts [Andrews et al., 2003]. However, recent decades have seen increased interest in the

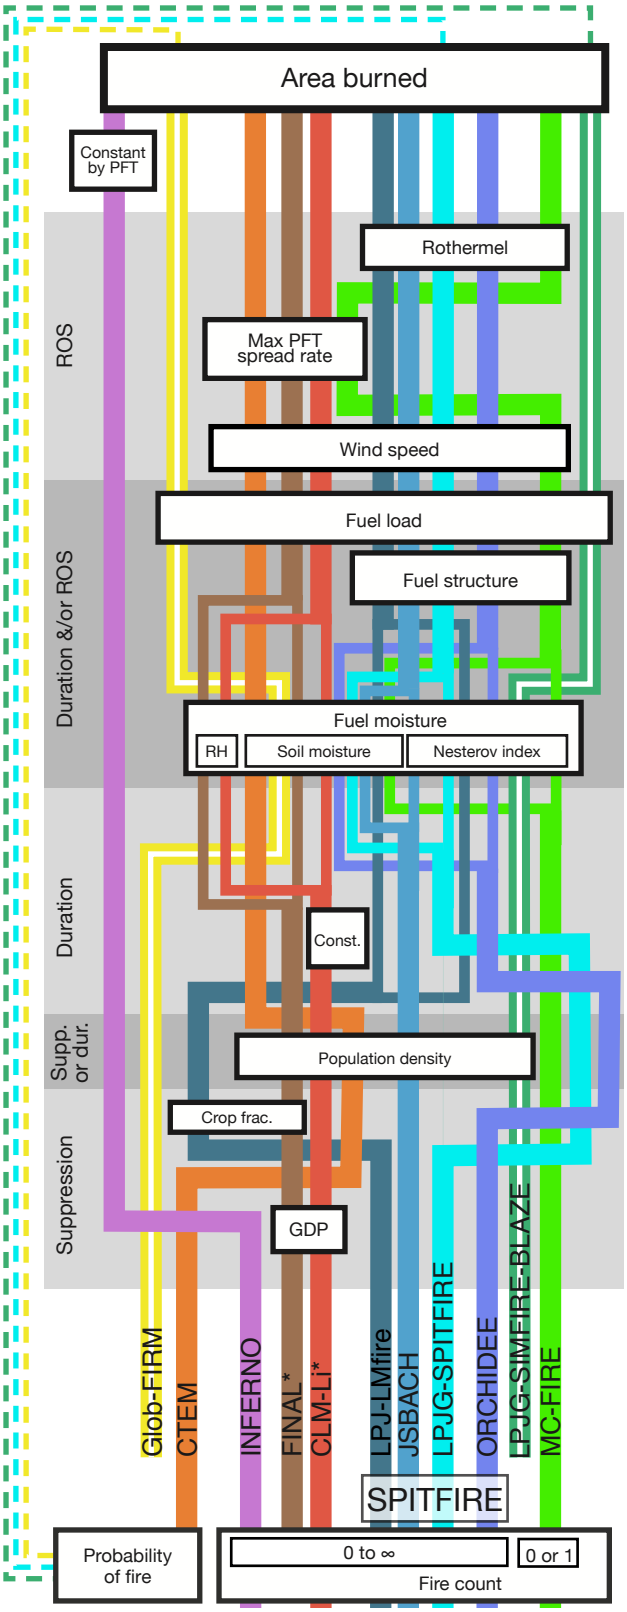


Figure 1.4: Processes leading from fire starts to BA in the FireMIP models. Source: Fig. 3 from Rabin et al. [2017], licensed under CC BY 3.0.

creation of larger-scale fire models designed for incorporation into Earth system models (ESMs) in order to better understand and predict our changing climate by improving our understanding of land-atmosphere feedbacks. Fire models have also been used to facilitate global studies of fire patterns before the advent of global remotely sensed fire data (i.e. the 1980s), by using climate data and fire-related biophysical variables to estimate past (and future) fire activity [Yang et al., 2014]. But fire models differ not only in their spatial scales but also in their temporal scales, with fire spread models used operationally with local weather operating on timesteps much shorter than 1 day [see e.g. Schumaker et al., 2022, for a recent example], while global models generally use large-scale average weather to estimate daily to annual fire risk, rather than fire spread.

Predicting the impact of changes in climate, resulting changes in vegetation, and human activities on fire regimes is complicated due to the numerous drivers and the feedbacks between them that need to be considered. Coupled fire-vegetation models can be used to estimate changes in large-scale fire regimes [see e.g. Kloster et al., 2012, Knorr et al., 2016] and explore how these may in turn affect regional vegetation since the aforementioned feedbacks can be represented in such models. Although several coupled fire-vegetation models have recently been developed, there are large differences in their predictions [Hantson et al., 2016, Forkel et al., 2019a, Teckentrup et al., 2019]. With the exception of LPJ-GUESS-SIMFIRE-BLAZE, recent decreases in global BA are generally not reproduced by fire models [Andela et al., 2017]. Furthermore, differences in model structure and parametrisation [Hantson et al., 2016] make it hard to discern whether correct simulations of modern fire regimes are indeed due to an accurate representation of the real underlying processes that control wildfires in nature. The model structure and the way that the underlying physics is treated [Perry, 1998] are also dependent on the spatial and temporal scales of interest. For example, small-scale topography and its effects on individual fires are unlikely to have an impact on larger scales (e.g. 100 km) [Archibald et al., 2009], while the effects of terrain heterogeneity as an agent of landscape fragmentation do have an impact [Archibald et al., 2009, Haas et al., 2022]. Similarly, the coarse-scale 0.5° grid cell analysis carried out by Haas et al. [2022] showed that mean monthly wind speed promotes fire spread but inhibits BA.

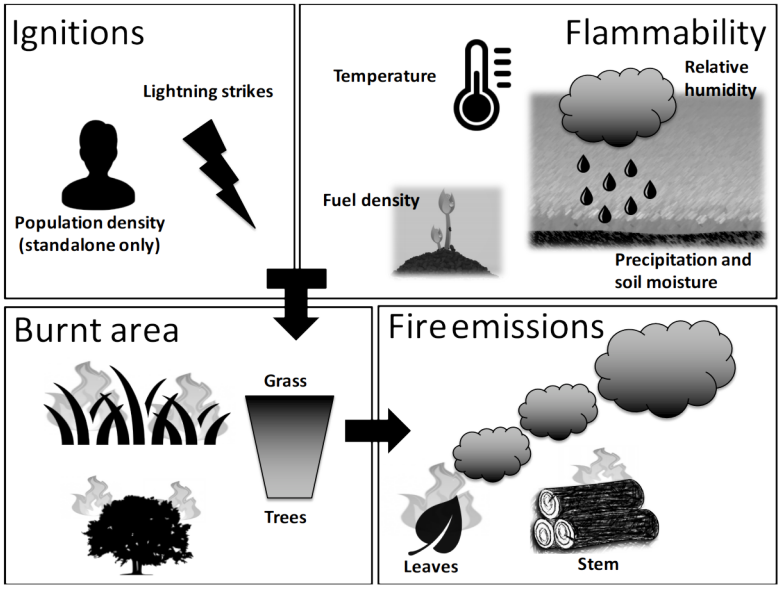


Figure 1.5: Underlying processes included in the INFERNO fire model. Multiple ignition sources, as well as biophysical variables together with the type of vegetation (PFTs), are seen to inform the amount of burnt area, which is then used to calculate emissions using conversion factors. Source: Fig. 1 from Mangeon et al. [2016b], licensed under CC BY 3.0.

This supports the earlier findings of Lasslop et al. [2015], who found no evidence for a positive global relationship between wind speed and BA beyond a relatively low wind speed threshold of roughly 2–3 $\mathrm{m\,s^{-1}}$. This may be due to the lack of human ignitions when wind speeds are high and conditions are conducive to fire spread to avoid escaped fires (see also [Magi et al., 2012]) or effects of wind on vegetation and fuel structure that could not be included in the analysis [Lasslop et al., 2015]. Understanding such discrepancies between observations and models, but also between models themselves (see Fig. 1.7), is the goal of the FireMIP project [Hantson et al., 2016, Rabin et al., 2017, Hantson et al., 2020].

1.5.1 The INFERNO model

The INteractive Fire and Emission algoRithm for Natural envirOnments (INFERNO) model employs simple parametrisations for predictions meant for application to decadal timescales or longer and was introduced by Mangeon et al. [2016b]. INFERNO was chosen for analysis in this thesis because it is an empirical fire model with a parsimonious approach to its driving variables. For example, it foregoes any representation of the influence of wind in favour of using a typical

BA for each plant functional type (PFT), as well as assuming that BA depends on the number of ignitions and a PFT-specific flammability (see Fig. 1.5). This not only agrees with the findings of Lasslop et al. [2015] and Haas et al. [2022] but also allows for a more straightforward adjustment of the relationships used within the model.

In contrast to the majority of FireMIP models (see Fig. 1.4), this parametrisation circumvents the use of wind (gusts) to model fire spread. At the large spatial scales of global ESMs, such as the $1.875° \times 1.25°$ resolution INFERNO was developed for, wind [Lasslop et al., 2015, Haas et al., 2022], topography [Archibald et al., 2009], and similar smaller-scale influences, like vegetation gap dynamics, are not expected to be the most relevant, and are therefore not explicitly modelled within INFERNO. This simplification also reduces computational requirements, facilitating coupling into an ESM which may be required to run over long periods. Despite its parsimonious approach, INFERNO performs as well as process-based models on the coarse scales investigated by Hantson et al. [2020].

Examples of empirical models besides INFERNO include the simple global fire model (SIMFIRE) introduced in Knorr et al. [2014], which uses a flexible parametrisation based on the logistic function to calculate fire frequency based on anthropogenic influence, fraction of absorbed photosynthetically active radiation (FAPAR), and the Nesterov index. The empirical GLOBal FIRe Model (Glob-FIRM) proposed by Thonicke et al. [2001], one of the first global fire models, uses a parametrisation based around the fire season length in a given fire year (based itself on fuel moisture, i.e. upper soil moisture, and the moisture of extinction) to estimate annual BA. Predating the widespread use of remotely sensed BA datasets, this approach also featured calibration using site-based observations.

1.5.1.1 Calculation of BA in INFERNO

The Joint UK Land Environment Simulator (JULES) is a community land surface model (LSM) based on the Met Office Surface Exchange Scheme (MOSES) [Best et al., 2011, Clark et al., 2011], able to represent a wide range of processes relating to vegetation or soil properties, for

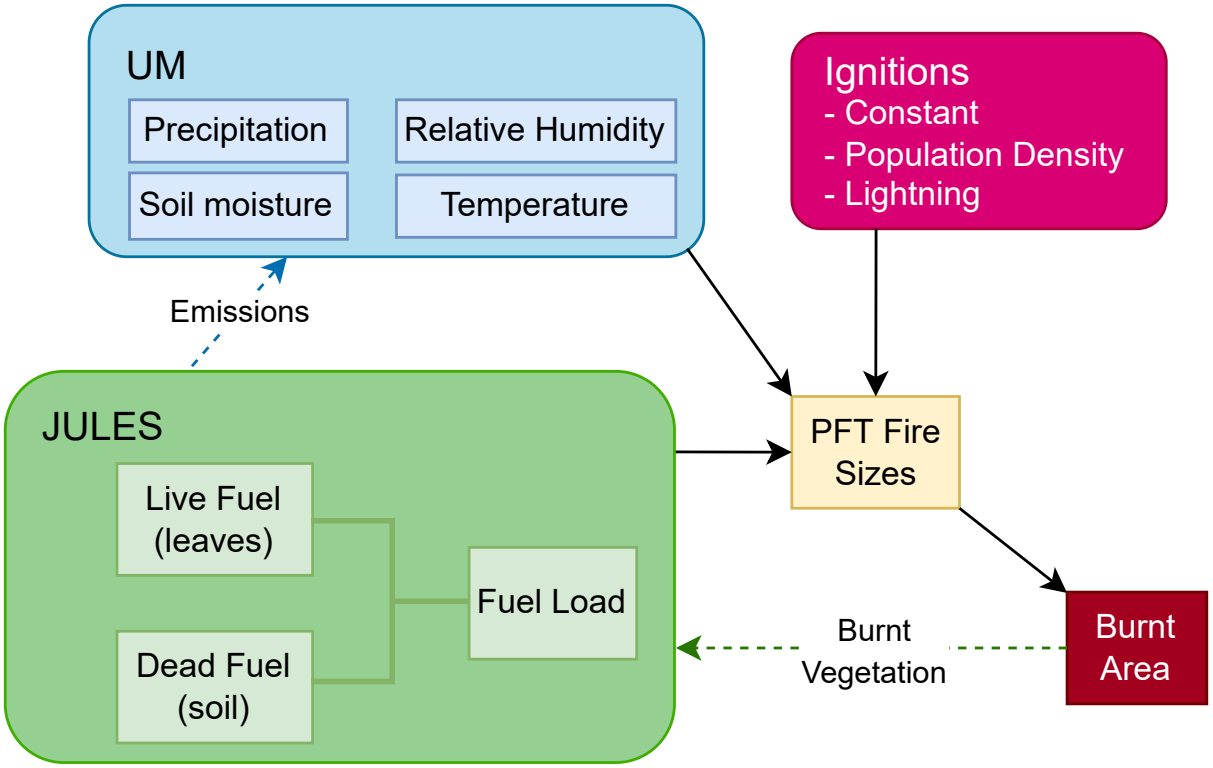


Figure 1.6: The fire-related linkages between different components of the UKESM. Flammability and fuel loads, along with the PFTs present in a grid cell and the number of ignitions, combine to yield BA. Soil carbon pools within the JULES land surface model are shown to inform fuel loads, with fire disturbance (BA) feeding back explicitly to JULES [Burton et al., 2019]. Fire emissions are also shown to impact the unified model (UM; atmospheric model) [Teixeira et al., 2021]. Lightning from the UKCA model is shown to influence ignitions. Adapted from Kelley et al. [2020] and [Mangeon et al., 2016b].

example. It can be forced by observations or coupled to an atmospheric model [Best et al., 2011] (e.g. within the UK Earth System Model (UKESM)). Vegetation properties are provided by JULES, and the coupled JULES-INFERNO model is run on individual grid cells. For each PFT in a grid cell, INFERNO determines BA using the product of the estimated number of fires and a typical BA per fire. As shown in Fig. 1.6, JULES and INFERNO are coupled by the influence that vegetation has on fuel loads and, conversely, the impact of fire disturbance on vegetation cover and soil carbon stocks [Burton et al., 2019]. Both live and dead fuel, meteorology, and ignitions influence BA, which then affects vegetation disturbance, soil carbon, and emissions.

The number of fires is the product of the number of ignitions and a combination of flammability and fuel load, both of which are PFT-specific. The number of ignitions is calculated for each grid cell and shared between PFTs according to one of three modes:

1. Constant, ubiquitous ignitions: $1.67\,\text{ignitions}\,\text{km}^{-2}\,\text{month}^{-1}$.
2. Constant human ignitions (and suppressions) at $1.5\,\text{ignitions}\,\text{km}^{-2}\,\text{month}^{-1}$ in addition to variable cloud–ground (CG) lightning strikes, every one of which is assumed to cause a fire.
3. Following Venevsky et al. [2002], in addition to variable CG lightning strikes as above, variable human influence may be prescribed using an external dataset, with effective ignitions dependent on POPD and HDI [Teixeira et al., 2021]. Potential anthropogenic ignitions increase with POPD and decrease with HDI. Finally, the fraction of fires that are not suppressed by humans decreases with both POPD (exponentially) and HDI (linearly).

$$F_{\text{PFT}} = \begin{cases} e^{*} e^{-2R} \text{FL}_{\text{PFT}} (1-\theta) & \text{for } \text{RH}_{\text{up}} < \text{RH} \\ e^{*} \dfrac{\text{RH} - \text{RH}_{\text{low}}}{\text{RH}_{\text{up}} - \text{RH}_{\text{low}}} e^{-2R} \text{FL}_{\text{PFT}} (1-\theta) & \text{for } \text{RH}_{\text{low}} < \text{RH} < \text{RH}_{\text{up}} \\ 0 & \text{for } \text{RH}_{\text{low}} > \text{RH} \end{cases} \tag{1.1}$$

Each PFT is assigned a fuel load (FL_{PFT}) and flammability (F_{PFT}, Eq. (1.1), from Mangeon et al. [2016b, Eq. (6)]) (see also Fig. 1.6). Flammability is determined by monthly averages of meteorological variables that are given to INFERNO via JULES: relative humidity (RH), precipitation (R in Eq. (1.1)), soil moisture (θ in Eq. (1.1)), and temperature (which determines the saturation vapour pressure, e^{*} in Eq. (1.1)). The fuel load depends on the amount of carbon contained in two soil carbon pools and above-ground leaf carbon, which are both simulated by JULES, with the amount of simulated leaf carbon being PFT-dependent. These decomposable plant material (DPM) and resistant plant material (RPM) carbon pools are shared amongst all PFTs, with the litter carbon flux (leaf, stem, and root carbon) from all JULES vegetation feeding into them in PFT-specific proportions [Clark et al., 2011]; for example, trees yield a higher proportion of resistant litter input than grasses. Both pools are intended to be used as flammable litter proxies [Burton et al., 2019]. The DPM pool is assumed to be 70 % accessible for consumption by fires (i.e. close to the surface).

A key simplification employed by the model is the usage of a prescribed constant PFT-specific BA per fire: $0.6\,\text{km}^2\,\text{fire}^{-1}$ is used for all tree PFTs, $1.4\,\text{km}^2\,\text{fire}^{-1}$ for all grass PFTs, and

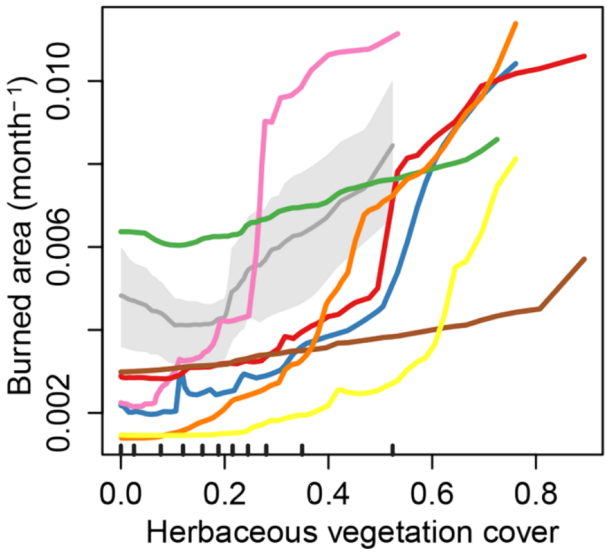


Figure 1.7: Global sensitivities of BA to herbaceous vegetation cover. INFERNO (JULES) is shown in brown. Five observed BA products are shown in grey (mean and range). Source: Fig. 4e from Forkel et al. [2019a], licensed under CC BY 4.0.

$1.2\,\mathrm{km}^2\,\mathrm{fire}^{-1}$ for all shrub PFTs.

Fuel build-up and the effect of antecedent conditions are represented by a combination of processes. The soil carbon pools accumulate carbon over time while also losing it via respiration. This accumulation of fuel is affected by antecedent vegetation disturbance, e.g. burning, as well as antecedent meteorological conditions and the modelled leaf and vegetation phenology. Furthermore, since INFERNO uses monthly averages of certain meteorological quantities, the antecedent conditions of these variables are used by definition. Precipitation, evapotranspiration, and runoff processes dictate the modelled soil moisture. The effect of antecedent precipitation on fuel moisture is represented using the top layer (up to 0.1 m depth) unfrozen soil moisture, where increasing fuel moisture leads to a decreasing amount of fire.

1.5.2 Model evaluation

Comparisons of FireMIP models have yielded important insights into the performance of state-of-the-art fire models. Andela et al. [2017] found that the models tend to struggle to predict the observed trend in BA, as well as the observed mean global BA. As stated previously, however, the uncertainty caused by the short observational record for satellite BA data (see [Forkel et al., 2019b]) currently prevents us from confidently stating which of the tested models performed

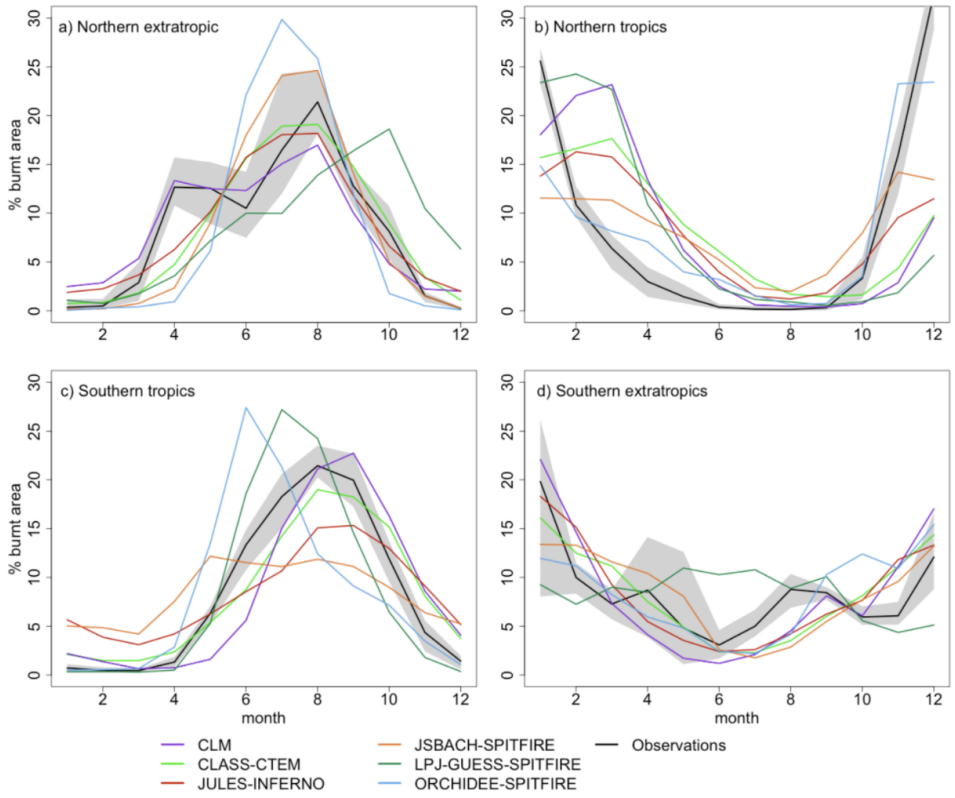


Figure 1.8: Seasonality of simulated and observed BA for 2002–2012. The black line indicates the mean of the remotely sensed BA datasets (GFED4, GFED4s, MCD45, FireCCI40, FireCCI51), while the grey shading indicates their min-max range. Source: Fig. 3 from Hantson et al. [2020], licensed under CC BY 4.0.

best in this test. However, the large discrepancies between them reinforce the need for further understanding of the underlying processes.

Forkel et al. [2019a] showed that FireMIP models captured the relationship with climate variables well, in contrast to the relationship with vegetation variables. Especially the dependence of BA on pre-season productivity was not represented well. This supports the findings of Hantson et al. [2020], who found that a better representation of the seasonal concentration of gross primary productivity (GPP) and vegetation patterns by the FireMIP models was correlated with better performance when predicting BA. This indicates that the seasonality of vegetation growth and senescence is important for the accurate modelling of fuel loads and, by extension, fire activity. Teckentrup et al. [2019] found that spatial patterns are captured well by the FireMIP models, but analysis of temporal trends since 1921 revealed large disagreements between the models. They attributed these differences largely to uncertainties related to the treatment of human impacts

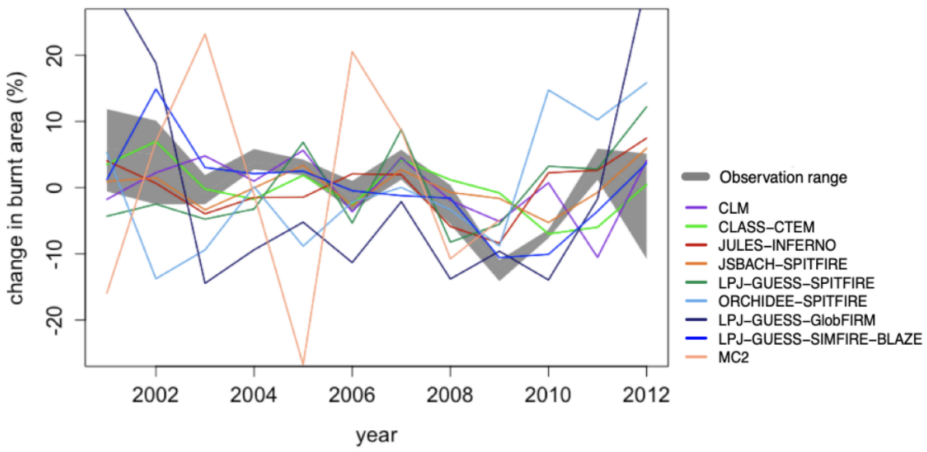


Figure 1.9: IAV of simulated and observed BA for 2001–2012. The grey shading indicates the min-max range of the remotely sensed datasets covering the entire time period (GFED4, GFED4s, MCD45, FireCCI51). Source: Fig. 4 from Hantson et al. [2020], licensed under CC BY 4.0.

(ignitions and suppression) and land use changes. While Hantson et al. [2020] also found that the general geographical pattern of BA is represented well by the FireMIP models (see Fig. 1.3), poor representation of fuel build-up by the models may have contributed to their comparatively poor reproduction of seasonality (see Fig. 1.8) and IAV (see Fig. 1.9) of BA.

Overall, Hantson et al. [2020] found that no single model outperformed all other models on all accounts and that empirical models such as INFERNO performed as well as the process-based models included in FireMIP. However, Jones et al. [2022] caution that empirical models often perform well because they are trained on observational data, with their ability to reliably extrapolate fire behaviour being unclear.

INFERNO has also been shown to exhibit an anomalous response to herbaceous vegetation cover, severely underestimating its influence on BA, as seen in Fig. 1.7 [Forkel et al., 2019a].

Simulations of contemporary fire emission totals by most FireMIP models were found to be within observational uncertainties [Li et al., 2019]. However, van Marle et al. [2017b] found large differences between the FireMIP models for simulated emissions between 1750 and 2015. As with BA [Hantson et al., 2020], Li et al. [2019] found that representation of IAV of fire emissions is lacking. They suggested that this, along with the discrepancy between modelled emission trends, may be due to a poor representation of human impacts in the models.

Beyond the systematic comparisons enabled by the FireMIP project, other studies have also evaluated fire models. Projections of future fire by models have been found to be subject to large uncertainties, with little agreement between the models [Settele and Scholes, 2014, Kloster and Lasslop, 2017].

1.6 Thesis structure

Chapter 2 introduces analytical methods used throughout the thesis for the modelling of BA and its drivers, such as generalised additive models (GAMs) and random forest (RF) regression models.

Chapter 3 describes the empirical wildfire analyses that have been published in Kuhn-Régnier et al. [2021b], albeit with additional results and methodological detail. The relationships between vegetation properties and BA are currently not understood as well as the relationships with climate [e.g. Forkel et al., 2019a, Kelley et al., 2019, Teckentrup et al., 2019, Hantson et al., 2020]. Therefore, both Forkel et al. [2019a] and Hantson et al. [2020] argued that an improved understanding of the relationship between vegetation properties and fuel build-up and their effects on fire regimes is needed to improve fire models. Using machine learning analyses of remote sensing datasets, my goal is to understand the fundamental drivers of BA. Particular emphasis is placed on the roles of antecedent vegetation productivity and aridity and how their impact relates to instantaneous conditions. Key to gaining such insights are machine learning methods such as RF regression models, which are more flexible than GLMs [e.g. Bistinas et al., 2014, Haas et al., 2022] and have previously shown promise as a tool for the analysis of wildfire drivers [Archibald et al., 2009, Aldersley et al., 2011, Forkel et al., 2019a]. They allow the representation and analysis of complex non-linear relationships between the driving variables and BA. Robust visualisations of these relationships are achieved using accumulated local effects (ALE). Interactions between the driving variables can also be quantified and visualised using RFs and ALEs. While statistical methods have previously been used to investigate fire drivers, including antecedent effects [e.g. Archibald et al., 2009, Forkel et al., 2017, Joshi and Sukumar, 2021, Haas et al., 2022], my

investigation aims to visualise the relationships between driving variables and BA on different antecedent timescales while also quantifying their importance for BA prediction at different antecedent timescales globally at a grid cell scale. This is enabled by correlation-tolerant analysis methods.

Chapter 4 details how the insights gained from Chapter 3 were used to develop a new flammability parametrisation for the INFERNO fire model. This parametrisation makes use of sigmoidal relationships similar to those used by Knorr et al. [2014] in order to enable the representation of a variety of relationships (e.g. linear or step-like) between the driving variables and flammability. To address the concerns that especially vegetation properties and fuel build-up are currently not represented adequately in global fire models [Forkel et al., 2019a, Hantson et al., 2020], particular emphasis is placed on re-parametrising the impact of antecedent vegetation, which, in addition to the expected climatic drivers, is found to be of particular importance in Chapter 3.

Chapter 5 then uses multiple sensitivity analyses to explore the underlying components of the parametrisation developed in the previous chapter. For example, the Sobol global sensitivity analysis (SA) methodology is used to quantify the impact different components have on either the phase or magnitude performance of the modelled BA. While simple fractional perturbations (i.e. +X%, -Y%) have previously been used to carry out SA for INFERNO in Mangeon [2017], Sobol SA yields more robust estimates. Furthermore, GAM models were fit to the modelled BA to verify that the sigmoidal relationships reasonably capture the underlying relationships. In a more fine-grained approach to the land cover grouping scheme variations used in Forkel et al. [2017], the ideal complexity (i.e. the number of parameters) of the parametrisation from Chapter 4 is explored by varying both the input data and the underlying relationships. Overall, the methods presented in this chapter serve to illustrate which parameters (and how many) and which input data of the new parametrisation are required for adequate prediction of BA while also revealing which aspects of the modelled BA are controlled by which parameters.

Finally, Chapter 6 provides a general discussion of the conclusions drawn from this work.

Links are made between the anticipated importance of different data and parameters from Chapter 3 based on observations and the eventual importance in a fire model context given the sensitivity analyses carried out in Chapter 5. It is also suggested how the approaches employed throughout may be used to analyse and improve other kinds of fire models, e.g. process-based models, in addition to how our general understanding of wildfire drivers might benefit from the insights gained in this thesis.

Chapter 2

Analysis Methods

Throughout the thesis, I used random forests (RFs) (Chapter 3) and generalised additive models (GAMs) (Chapter 5) for statistical modelling of burnt area (BA). Accumulated local effects (ALE) were used to visualise the underlying relationships fitted by the RFs in Chapter 3.

2.1 Emergent vs underlying relationships

It is important to consider the multiple factors that influence fire together and to account for their interactions in order to distinguish underlying from emergent relationships [Yang et al., 2014, Bistinas et al., 2014, Archibald et al., 2018]. Several approaches have been used, ranging from generalised linear models (GLMs) [e.g. Bistinas et al., 2014, Haas et al., 2022] to machine learning approaches [e.g. Archibald et al., 2009, Aldersley et al., 2011, Forkel et al., 2019a, Joshi and Sukumar, 2021].

Using GLMs, Bistinas et al. [2014] demonstrated the significant differences between emergent and underlying (or 'causal') relationships between drivers of wildfires and BA. An example of the discrepancy between emergent and underlying relationships is the relationship between temperature and BA. While the underlying relationship between the two variables is positive and monotonic as expected (higher temperatures result in higher BA), the emergent relationship, observed e.g. when plotting observations of temperature against observed BA, is unimodal with a single peak at intermediate temperatures. The latter relationship would suggest a decrease in BA

with increasing temperatures at high temperatures. However, this is an artefact of other factors, e.g. the lack of vegetation in regions with extreme temperatures. Underlying relationships can be determined by taking into account all other relevant variables during statistical modelling (see Section 2.3). Thus identified, such relationships can then augment existing ecological knowledge to help infer underlying processes that control fires.

2.2 Generalised additive models

Generalised additive models (GAMs) can be seen as an extension of GLMs. Compared to simple linear models, GLMs are more flexible by being able to take into account different error distributions and transformations of the mean response [Nelder and Wedderburn, 1972]. Through the use of a link function, GLMs allow separate modelling of the linearity and variance, enabling a more flexible treatment than simple variable transformations. To further increase the flexibility of the model compared to GLMs, GAMs use an additive combination of smoothing functions (e.g. polynomials or non-parametric functions; f_n in Eq. (2.1)) [Hastie and Tibshirani, 1999]. Because of this, GAMs can model non-linear and non-parametric relationships between the predictor and target variables. The GAM model can be written as follows:

$$g(E(Y)) = g(\mu) = \eta = \beta_0 + f_1(x_1) + f_2(x_2) + \cdots + f_m(x_m), \tag{2.1}$$

where $\mu = E(Y)$ is the expected value of the response variable Y, g denotes the link function, and β_0 denotes the intercept.

2.3 Random Forests

Random forest (RF) regression is an ensemble learning approach in which multiple decision trees are constructed using a randomly sampled subset of training observations (see Fig. 2.1) [Breiman, 2001]. The output of the final model is the average result of all individual decision trees. RF regression is highly suited to investigating the emergent controls on fire because it is able to learn

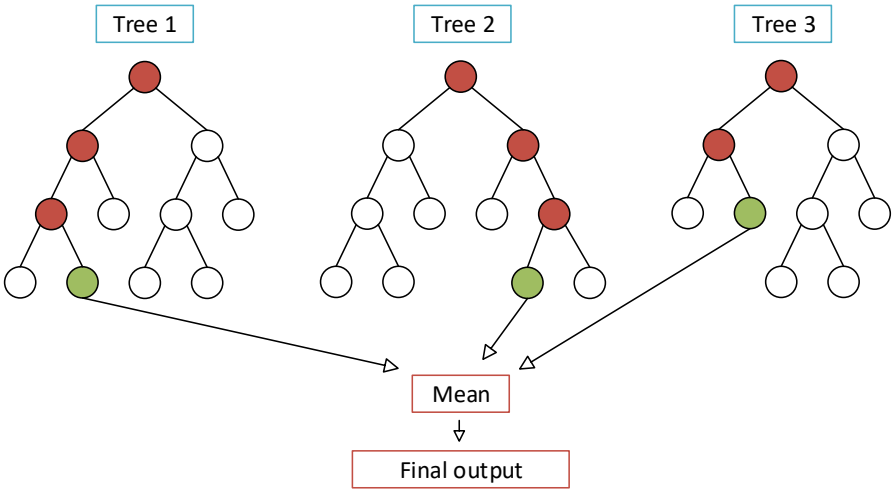


Figure 2.1: The random forest algorithm. Multiple trees are fit using potentially different bagged samples of the training data. For regression, the final prediction is the mean of the individual regression tree predictions. For each regression tree, samples are propagated based on the value of several variables. At every tree node, a variable is compared to some threshold to decide which branch is followed (red nodes) and which is discarded (white nodes). Finally, the lowest tree nodes (green) denote the output value for each tree given the sequence of previous comparisons.

non-linear relationships in high-dimensional space [Archibald et al., 2009]. By averaging over multiple decision trees, RFs also aim to mitigate overfitting [Breiman, 2001]. Compared to GAMs, the possible relationships between the predictor and target variables are more flexible, with arbitrary interactions between variables also being fitted. While other machine learning (ML) approaches have been used to study fire, e.g. multilayer neural networks [Joshi and Sukumar, 2021], I chose RFs for my analysis since they have previously been shown to work well for studies of fire (see e.g. [Archibald et al., 2009, Aldersley et al., 2011, Forkel et al., 2019a]). They are also more interpretable than other approaches like neural networks. Furthermore, an efficient method for SHapley Additive exPlanations (SHAP) value calculation has been formulated for such tree-based methods [Lundberg et al., 2020].

Fitting RFs is a supervised problem, meaning that both the predictor and target variables are used to estimate ('fit') relationships between the two. The exact mechanism used for the model fitting procedure is dependent on several parameters, known as 'hyperparameters', the optimal values of which depend on the dataset and performance metric (e.g. R^2) used [Probst et al., 2019]. Examples of hyperparameters include the number of estimators (`n_estimators`),

which determines the number of trees whose predictions are averaged and the maximum depth (`max_depth`), which limits the number of split levels. Generally, increasing the number of estimators (trees) leads to a performance improvement, albeit at the expense of increasing the computational cost [Probst et al., 2019]. Reducing the maximum depth 'prunes' the trees, reducing the number of tree nodes that only depend on a few samples. This can lead to improved predictive performance, especially in the presence of noise [Segal, 2004, Probst et al., 2019]. Intuitively, pruning may reduce overfitting because reducing the number of nodes is akin to disregarding the finest details, which are likely to be more affected by noise as opposed to being representative of a genuine underlying relationship.

2.3.1 N-fold cross-validation

Hyperparameter estimation is commonly carried out using a technique called N-fold cross-validation (CV), whereby N distinct non-overlapping slices of data are separated into training and validation datasets [Fox et al., 2017, Roberts et al., 2017]. $N-1$ of the data slices are used for training, with the model performance then being tested (e.g. using the R^2 metric) on the remaining data slice. In the same fashion, all data slices are alternately used for testing, with the final result being the average model performance (e.g. R^2) across all test slices. Compared to regular validation, where a certain fraction of data is (perhaps randomly) segregated into fixed training and validation datasets (i.e. a single step in N-fold CV), N-fold CV evaluates model performance on more varied samples without requiring more data. This is effective since model training is reset between runs (the model is retrained 'from scratch' for every new combination of training dataset slices).

2.4 Machine learning model validation

Model performance on the validation dataset, which has not been used for training (including any hyperparameter tuning), is an indication of the expected performance of the model on unseen data (i.e. model generalisability). For the RF model, the out-of-bag (OOB) score also provides

an alternative way of evaluating model performance [Breiman, 2001, Segal, 2004, Fox et al., 2017, Probst et al., 2019]. Individual trees of an RF model are constructed using a bootstrap sample (random sampling with replacement). Owing to this sampling scheme, around a third of all data are not contained within this sample and are referred to as the tree's OOB data. Predictions on OOB samples can then be used to derive the OOB score [Breiman, 2001]. Therefore, the OOB R^2 for the training dataset can be used to estimate how the model will perform on unseen data [Segal, 2004, Fox et al., 2017, Probst et al., 2019]. However, since these data are still part of the training dataset, this may skew the expected performance. As such, the R^2 score for the validation dataset can be used to provide an independent, alternative measure of the model's generalisability.

2.4.1 Spatial autocorrelation

Use of the validation dataset R^2 is only valid if there is no autocorrelation between the samples. In the presence of spatial autocorrelation, the data are no longer independent [Archibald et al., 2009], as samples that are spatially close to each other tend to be similar [Ploton et al., 2020]. Spatial autocorrelation may impact model fitting and validation [Ploton et al., 2020] since the validation samples are not independent of the training samples used to fit the model.

Following the procedure outlined by Ploton et al. [2020], the degree of spatial autocorrelation can be estimated using a variogram of the global target variables (BA). This is then followed by a buffered leave-one-out cross-validation (B-LOO CV) procedure to determine the extent to which autocorrelation may influence the amount of potential overfitting of the model used. This is carried out as follows, where $r_{\max}$ is inferred from the variogram by noting at which separation the semivariance stops increasing, and N_t is chosen such that the number of potential training samples is guaranteed to be equal to or above N_t for all $r \leq r_{\max}$:

1. Randomly choose a single (test) location. The samples (e.g. monthly samples) at this test location constitute the test set.
2. Exclude samples from the potential training set in a circular region of radius r pixels

around the test location, such that no potential training sample is closer than r pixels to the test location. This limits the influence of spatial autocorrelation.

3. Randomly choose N_t training samples from all remaining potential training samples.
4. Using a model trained on the above training samples, predict BA for the test location.
5. Increment r and repeat steps 1–4 until r has reached $r_{\max}$.

This process can be repeated many times for each of the chosen radii by changing the random seed associated with the random test location choice.

2.5 Visualising underlying relationships using ALE plots

Black-box ML models, including reasonably complex RFs that cannot be visualised feasibly in terms of their underlying regression trees, are powerful regression tools due to their ability to capture arbitrarily complicated relationships between input and output data. However, this complexity goes hand in hand with the inability to easily infer the average response of these models to variations in the input data, i.e. the underlying relationships fit by the model. To enable the visualisation of such underlying relationships, partial dependence plots (PDPs) are widely used due to their intuitive interpretation and implementation [Zhao and Hastie, 2019, Molnar, 2020], including for RF wildfire predictions [Forkel et al., 2019a]. They are computed by fitting the model to every possible combination of independent input variables while keeping one of them constant. The resulting predictions are then averaged in order to remove the effects of all but the variable that was kept constant. Repeating this process while keeping the variable in question constant at a variety of values isolates the influence of the chosen variable across the chosen range [Friedman, 2001].

ALE plots represent a more statistically rigorous and robust alternative to PDPs [Molnar, 2020, Apley and Zhu, 2020]. They can be interpreted in much the same way as PDPs while being statistically favourable for correlated input variables. This is because, in the presence of correlations, computing PDPs forces the model under investigation to extrapolate beyond the distribution of input variables present in the training set, leading to misleading results [Hooker,

2007, Molnar, 2020, Apley and Zhu, 2020]. Especially for models that are poor at extrapolating, such as the RF model, this may result in unpredictable and, therefore, undesirable behaviour. Unlike e.g. GLM models, which fit continuous relationships between predictor variables and the target variable that is to be predicted, RF models partition the input space into a set of categories, each associated with a certain target value that is within the values seen in the training set. A direct corollary of this approach is the inability to predict target values that are higher or lower than the values seen during training since all predictions must be associated with one of the categories learnt during the RF model training. The ALE technique circumvents this problem by only utilising a limited number of predictor variable combinations in the vicinity of existing training samples. This has the convenient side-effect of reduced computational runtime since fewer predictions need to be made. These runtime reductions are especially pronounced for the 2D ALE plots explained below.

2.5.1 Visualising inhomogeneities

Underlying inhomogeneities appear when the relationships fit by a model differ between regions (and times), as is expected to happen when comparing biomes, for example. These are traditionally revealed using individual conditional expectation (ICE) plots that isolate the modelled response to single features for a given instance [Goldstein et al., 2015] (e.g. a single location and month). However, ICE plots suffer from the same problems as PDPs since they are computed in much the same way (PDPs are mathematically equivalent to an average over ICEs). Contrary to ICE plots, the construction of ALE plots for single samples is not possible by design. However, the ALE technique can be used to examine inhomogeneities by subsampling the dataset prior to ALE plotting, thereby allowing the visualisation of underlying relationships for a subset of locations and times. Causes of inhomogeneities can be explored using 2D ALE plots, which show the combined effect of two features on the predictions.

Chapter 3

Empirical Analysis of Burnt Area Drivers

The work presented in this chapter has been published in part in the following article:

> Alexander Kuhn-Régnier, Apostolos Voulgarakis, Peer Nowack, Matthias Forkel, I. Colin Prentice, and Sandy P. Harrison. The importance of antecedent vegetation and drought conditions as global drivers of burnt area. *Biogeosciences*, 18(12):3861–79. June 2021. ISSN 1726-4170. doi: 10.5194/bg-18-3861-2021.

Author contributions to the work above are as follows:

> Conceptualisation: AKR, AV, ICP, SPH; Methodology: AKR; Software: AKR; Validation: AKR; Formal Analysis: AKR; Investigation: AKR, SPH; Data Curation: AKR, MF; Writing – Original Draft: AKR; Writing – Review & Editing: AKR, AV, PN, MF, ICP, SPH; Visualisation: AKR; Supervision: AV, PN, ICP, SPH.

Additional contributions presented here include Fig. 3.3 in Section 3.2.4.

Abstract

The seasonal and longer-term dynamics of fuel accumulation affect fire seasonality and the occurrence of extreme wildfires. Failure to account for their influence may help to explain why state-of-the-art fire models do not simulate the length and timing of the fire season or inter-annual variability in burnt area well. I investigated the impact of accounting for different timescales of fuel production and accumulation on burnt area using a suite of random forest regression models that included the immediate impact of climate, vegetation, and human influences in a given month and tested the impact of various combinations of antecedent conditions in four productivity-related vegetation indices and in antecedent moisture conditions. Analyses were conducted for the period from 2010 to 2015 inclusive. The inclusion of antecedent vegetation conditions representing fuel build-up led to an improvement of the global climatological out-of-sample R^2 from 0.579 to 0.701, but the inclusion of antecedent vegetation conditions on timescales $\geq$1 yr had no impact on the simulated burnt area. Current moisture levels were the dominant influence on fuel drying. Additionally, antecedent moisture levels were important for fuel build-up. The models also enabled the visualisation of interactions between variables, such as the importance of antecedent productivity coupled with instantaneous drying. The length of the period which needs to be considered varies across biomes; fuel-limited regions are sensitive to antecedent conditions that determine fuel build-up over longer time periods ($\sim$4 months), while moisture-limited regions are more sensitive to current conditions that regulate fuel drying.

3.1 Introduction

Studies have pinpointed the relationship between simulated vegetation properties and burnt area (BA) as a cause for concern [e.g. Forkel et al., 2019a, Kelley et al., 2019, Teckentrup et al., 2019, Hantson et al., 2020]. Forkel et al. [2019a] analysed satellite data to show that while state-of-the-art fire–vegetation models reproduce the emergent relationships with climatic variables, they do not correctly represent the relationship between vegetation and BA. Hantson et al. [2020] highlighted the need for an improved understanding of the vegetation drivers of fire season length and the inter-annual variability (IAV) of BA. Both Forkel et al. [2019a] and Hantson et al. [2020] argued for a better understanding of how vegetation properties control fuel build-up and, therefore, fire occurrence and intensity.

Several regional and global studies have indicated the importance of antecedent fuel build-up for BA. For example, links between fire activity and antecedent productivity have been found in South Africa [Van Wilgen et al., 2000], Australia [Griffin et al., 1983, Jenkins et al., 2020], grass and shrublands in the western United States [Swetnam and Betancourt, 1998, Westerling et al., 2003, Littell et al., 2009], and southern Africa [Archibald et al., 2009]. Global studies have identified similar relationships (a positive relationship between pre-season productivity and fire activity in the following dry season) in some dry areas. By studying the correlation between growing period (i.e. antecedent) soil moisture and fire activity, Krawchuk and Moritz [2011] found that fire activity in dry regions is related to antecedent productivity. Similarly, van der Werf et al. [2008] found a similar relationship for arid ecosystems (e.g. northern Australia), where antecedent wet conditions coupled with high drying rates in the month of the fire were found to be important. Some studies have also identified northern Australia to obey this relationship [Randerson et al., 2005, Spessa et al., 2005]. Alvarado et al. [2020] used 6-month and 24-month antecedent rainfall to demonstrate fuel moisture and fuel build-up limitations, respectively, across tropical savannahs, also finding a strong response of BA to antecedent precipitation in arid African and Australian regions. In another recent global analysis, O et al. [2020] found that for arid regions, wet anomalies (soil moisture) led to increased fire later in the year by increasing fuel loads and

biomass, supporting the regional findings of e.g. Krueger et al. [2016] for Oklahoma in the US. Thus, it is clear that a better understanding of the timescales of fuel accumulation, the interaction between biophysical drivers and fuel build-up, and the effects of antecedent weather conditions on both fuel loads and fuel drying is needed in order to improve BA predictions.

Here, I quantify the roles that antecedent vegetation productivity and aridity play relative to instantaneous conditions, the critical number of months that are most important for each, the shape of their relationships to BA, and the interactions between them. Since other climate factors, ignitions, and human activities also influence BA, I necessarily include these factors in my analysis. The use of a machine learning approach enables the identification of non-linear relationships and interactions between the drivers. This is then combined with analysis and visualisation techniques that provide insights into the modelled relationships while mitigating the effects of correlations among variables. Such insights include the effect of a particular driver on BA and the interactions between pairs of drivers. While other studies have used statistical methods to explore fire drivers, including the effect of antecedent productivity [e.g. Archibald et al., 2009, Forkel et al., 2017, Joshi and Sukumar, 2021, Haas et al., 2022], my study explores the global relationships between antecedent conditions (fuel load and drying) and fire in more detail. Compared to previous regional investigations of the (relative) importance of antecedent variables [e.g. Bessie and Johnson, 1995], I aim to quantify this on a global scale.

3.2 Methods

3.2.1 Data

The predictor and BA datasets are available for different but overlapping time periods (Table 3.1). I pre-processed each dataset separately and conducted random forest (RF) analyses based on the common period from January 2010 to April 2015 using the monthly climatological average of each time series. Monthly fractional BA for this period was obtained from the fourth-generation global fire emissions database (GFED4) dataset [Giglio et al., 2013] (data were retrieved from `https://www.globalfiredata.org/data.html`). To investigate the impact of temporal trends,

a longer time period from November 2000 to December 2019 was also considered in a monthly RF analysis using a subset of significant variables which were available for this extended time period. The MODIS MCD64CMQ [Giglio et al., 2018] BA dataset was used for this analysis to avoid the temporal limitation of the GFED4 dataset (see Table 3.1). Note that GFED4 is based on MODIS BA.

Diurnal temperature range (DTR), maximum temperature (MaxT), dry-day period (DD), and soil moisture are important climate factors influencing BA [Archibald et al., 2009, Bistinas et al., 2014, Forkel et al., 2017, Abatzoglou et al., 2018, Forkel et al., 2019a, Kelley et al., 2019] and are therefore considered as predictors in my analyses. The DTR was calculated by taking the monthly average of the difference between the daily maximum and minimum ECMWF Reanalysis 5th Generation (ERA5) [Copernicus Climate Change Service (C3S), 2017] 2 m temperatures.

The dry-day period was defined as the longest contiguous period of ERA5 mean daily precipitation below 0.1 mm day^{-1} [wetting rainfall; Harris et al., 2014, Jolly et al., 2015] within each month. Contiguous dry-day periods were concatenated across months before determining the longest dry-day period within a given month. For example, consider a 30-day long month with a 10-day long dry-day period at the beginning of the month, followed by a wetting precipitation event on day 11, and then a dry-day period for the following 19 days. This month has a dry-day period of 19 days. However, if the previous month were to terminate in a 10-day long dry-day period, these 10 days would be added to the initial 10-day dry-day period of the current month, thereby making this combined dry-day period the longest. Assuming that this previous month does not have any dry-day periods apart from the last 10-day-long dry-day period, this would make up the longest dry-day period in this month, regardless of the fact that these 10 days were added to the dry-day period in the following month.

Soil moisture was taken from the Copernicus soil water index (SWI) dataset [Wagner et al., 1999, Albergel et al., 2008]. The WWLLN Global Lightning Climatology (WGLC) dataset [Kaplan and Lau, 2019] was used to provide counts of monthly lightning strikes. It is based on the World Wide Lightning Location Network (WWLLN) dataset, which mainly detects cloud–ground (CG)

Table 3.1: Characteristics of the datasets. End times, as applicable to the processed data, are indicated in brackets. Rows are coloured to aid readability.

Variable	Abbreviation	Dataset	Start	End	Time	Reference
burnt area	BA	GFED4	06-1995	12-2016	monthly	Giglio et al. [2013]
diurnal temperature range	DTR	ERA5	1950	present (12-2020)	monthly	Copernicus Climate Change Service (C3S) [2017]
maximum temperature	MaxT	ERA5	1950	present (12-2020)	monthly	Copernicus Climate Change Service (C3S) [2017]
dry-day period	DD	ERA5	1950	present (11-2020)	monthly	Copernicus Climate Change Service (C3S) [2017]
soil moisture	SWI	Copernicus SWI	01-2007	11-2018	monthly	Wagner et al. [1999], Albergel et al. [2008]
lightning	Lightning	WGLC Lightning	01-2010	12-2018	monthly	Kaplan and Lau [2019]
above-ground tree biomass	AGB	Tropical: Avitabile, Northern: Thurner	static	static	static	Thurner et al. [2014], Avitabile et al. [2016]
land cover (fractional cover per grid cell)	CROP SHRUB TREE HERB	ESA CCI Land Cover	1992	present (2019)	yearly	Li et al. [2018]
solar-induced fluorescence	SIF	GlobFluo SIF	01-2007	04-2015	monthly	Köhler et al. [2015]
vegetation optical depth	VOD	VODCA (Ku-band)	12-1997	12-2018	monthly	Moesinger et al. [2019]
fraction of absorbed photosynthetically active radiation	FAPAR	MOD15A2H	02-2000	present (03-2021)	monthly	Myneni et al. [2015]
leaf area index	LAI	MOD15A2H	02-2000	present (11-2018)	monthly	Myneni et al. [2015]
population density	POPD	HYDE 3.2 (updated)	2000	present (2020)	yearly	Klein Goldewijk [2017], Kees Klein Goldewijk; personal communication (February 2021)
MCD64CMQ burnt area	MCD64 BA	MCD64CMQ	11-2000	present (06-2020)	monthly	Giglio et al. [2018]

strikes [Rodger et al., 2004, Abarca et al., 2010], as opposed to Lightning Imaging Sensor (LIS) lightning data [Bürgesser, 2017].

I included several alternative representations of land cover, including above-ground biomass (AGB) and the fractional cover of trees (TREE), shrubs (SHRUB), herbaceous vegetation (HERB), and crops (CROP), in my predictor set. Mean tree AGB was obtained by mosaicking AGB datasets for the tropics [Avitabile et al., 2016, 1 km resolution] and northern forests [Thurner et al., 2014, 0.01° resolution] using the mean after resampling each to a common spatial resolution of 0.25°. Yearly land cover values were obtained from the ESA Climate Change Initiative (CCI) Land Cover dataset [Li et al., 2018]. Land cover types were converted to fractional cover according to Poulter et al. [2015] using the conversion table from Forkel et al. [2017]. Global population density (POPD) from an updated version of the History database of the Global Environment (HYDE) 3.2 dataset [Klein Goldewijk, 2017, Kees Klein Goldewijk; personal communication, February 2021] was used as a measure of human influence on vegetation and fire regimes, as comparable indices (e.g. the night light development index), were shown to perform similarly [Forkel et al., 2017] and it is already widely used, as well as being easily interpretable. Regional proxies of human influence, like road density or housing density, are also unlikely to generalise to global scales and across time periods [Syphard et al., 2007]. Despite their yearly or even static timescales, the above datasets are valuable to my analysis because they provide additional information when combined with meteorological variables that vary on monthly timescales [Kelley et al., 2019].

Field data on fuel loads are sparse, and the only global dataset [Pettinari and Chuvieco, 2016] is based on extrapolating scattered field measurements by biome. I therefore used four remotely sensed vegetation properties related to total biomass or leaf cover that could be regarded as fuel load indices in my predictor set: solar-induced fluorescence (SIF), vegetation optical depth (VOD), the fraction of absorbed photosynthetically active radiation (FAPAR), and leaf area index (LAI). All four properties have previously been used as productivity indices [e.g. Ogutu et al., 2014, Teubner et al., 2018, Ryu et al., 2019, Mohammed et al., 2019] and I use all

four because it is uncertain which would be most closely related to fuel loads. There are also trade-offs involved in choosing between these remotely sensed properties. For example, at high VOD values above 1.2, small-scale canopy changes can no longer be detected [Owe et al., 2001, van Marle et al., 2016]. Monthly SIF was obtained from the GlobFluo SIF dataset [Köhler et al., 2015]. Ku-band VOD was obtained from the VOD Climate Archive (VODCA) dataset [Moesinger et al., 2020]. FAPAR and LAI were obtained from the MOD15A2H dataset [Myneni et al., 2015]. I used data from January 2008 to April 2015, covered by all four datasets, to derive values (including up to 2 years antecedent values) for the climatological analyses of the period from January 2010 to April 2015. Similarly, relevant data from February 2000 to December 2019 were pre-processed for the monthly analysis of the period from November 2000 to December 2019 (including up to 9 months antecedent values). The varying time periods required for data pre-processing for these two analyses result from the choice of antecedent variables used, as explained in Sections 3.2.1.1 and 3.2.3.

3.2.1.1 Antecedent predictor variables

The influence of antecedent conditions that might affect fuel loads or fuel dryness, specifically vegetation properties and DD, on BA was investigated using antecedent FAPAR, LAI, VOD, SIF, and DD data from up to two years before any given month (1M, 3M, 6M, 9M, 12M, 18M, 24M, where M denotes months). The large temporal autocorrelation between predictor variables could impede the visual interpretation of the impacts of antecedent periods $\geq$1 yr. Thus, anomalies were computed by subtracting the seasonal cycle relative to the designated month, resulting in the following transformations:

$$(\mathrm{X\ 12M}) - (\mathrm{X\ 0M}) \rightarrow \mathrm{X}\ \Delta 12\mathrm{M},$$

$$(\mathrm{X\ 18M}) - (\mathrm{X\ 6M}) \rightarrow \mathrm{X}\ \Delta 18\mathrm{M},$$

$$\text{and } (\mathrm{X\ 24M}) - (\mathrm{X\ 0M}) \rightarrow \mathrm{X}\ \Delta 24\mathrm{M},$$

where X ∈ {FAPAR, LAI, VOD, SIF, DD} and X 0M refers to the variable X in the current month. For example, the 12-month antecedent X 12M was transformed by subtracting the instantaneous (month 0) value of X, thereby yielding the anomaly in X, X Δ12M, which may be easier to interpret.

3.2.2 Data processing

3.2.2.1 Gap filling

There are gaps in the SWI, FAPAR, LAI, SIF, and VOD datasets in winter months at latitudes above ∼60°N, and in the austral winter for southern South America, due to high solar zenith angles for FAPAR, LAI, and SIF and because of snow cover and frozen soil for SWI and VOD [see e.g. Moesinger et al., 2020]. There are also sporadic missing values in these datasets caused by e.g. cloud cover. Simple exclusion of times lacking data is not feasible when examining antecedent conditions. Thus, data gaps were filled using a two-step approach, as in Forkel et al. [2017], in order to allow the analysis of summer months at the affected locations. This approach differentiates between two gap types based on the amount of missing information for a specific month at each location.

First, 'persistent' gaps, defined as months for which 50 % or more of the observations across all years are missing, were filled using the minimum value observed at that location for the given predictor variable. I assume that this indicates missing data during the winter since other causes for data gaps (e.g. cloud cover) are predominantly 'transient'. For example, if a certain grid cell was missing data for more than 50 % of all Decembers in the record, these gaps in December would be treated as persistent and therefore filled using minima.

Second, the remaining transient gaps were filled using season-trend regression models with four harmonic terms ($k = 4$) and without breakpoints. These models were fitted using ordinary least squares regression to the entire time series obtained during the first step, using data from January 2008 to April 2015 for the climatological analyses or February 2000 to December 2019 for the monthly analysis. Cloud cover, which also affects detection in tropical and subtropical

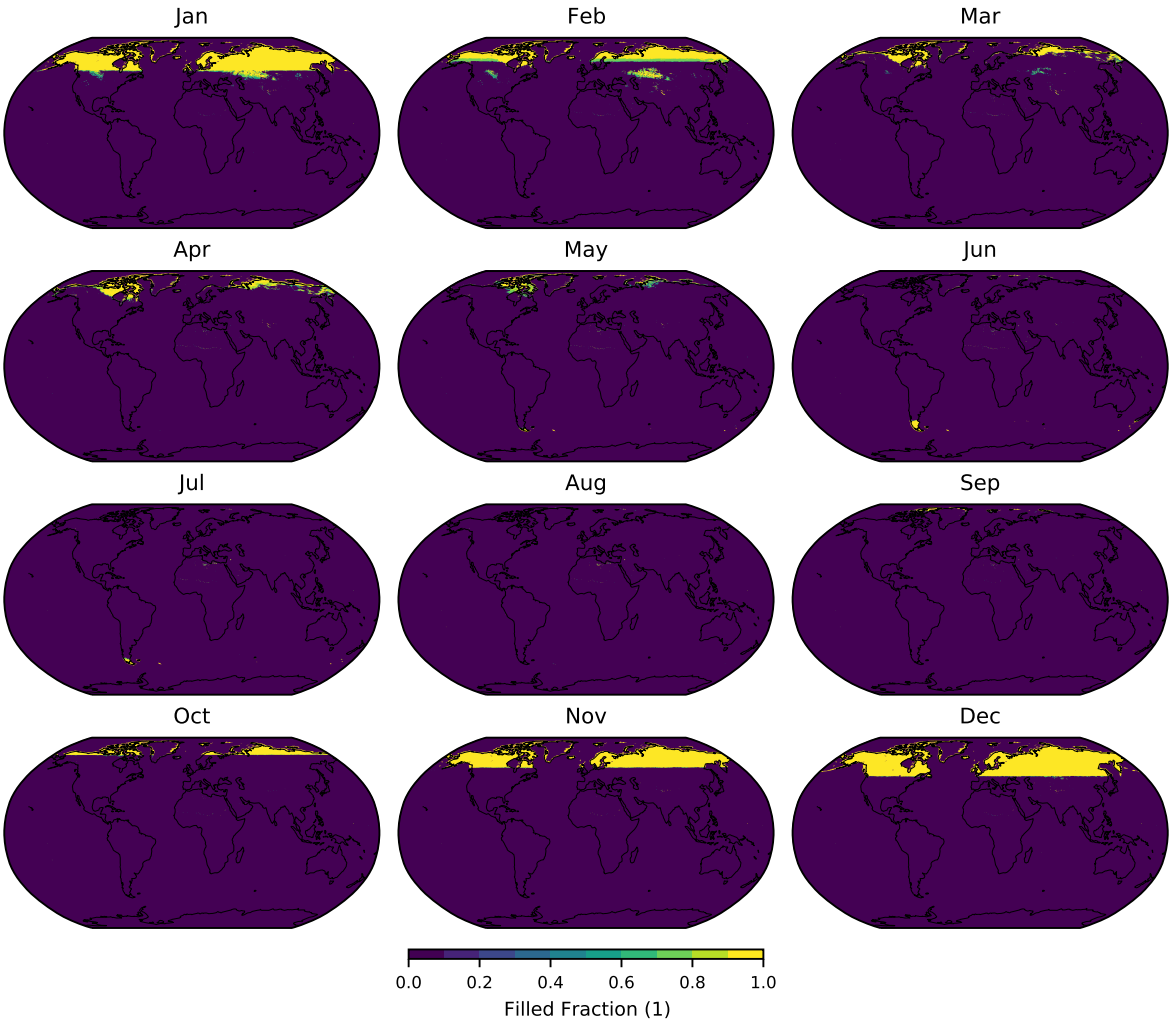


Figure 3.1: The fraction of filled samples for FAPAR (January 2008 to April 2015) at a given location for each month, with yellow indicating that all occurrences of a given month at a given location were filled and purple indicating that no filling was done. Filling is carried out mainly in winter in northern latitudes.

regions, is usually transient and therefore filled using the regression models. For the climatological analyses with 88 months in total, locations where no observations were available for > 52 months (regardless of whether such data gaps always occurred in the same month, as for persistent gaps, or at any point throughout the year) were discarded in a trade-off between data quality and geographic extent. For the monthly analysis, locations were discarded given > 138 unavailable months out of 239 months.

An alternative temporal nearest-neighbour gap-filling approach yielded very similar results (Fig. 3.10b). This simple approach used for the eventual ALL_NN model processes time series at a given location, filling gaps by using the temporally closest available sample(s) at that location.

Of the two approaches, I used the season-trend model with minima filling because it represents a more physical solution; it is based on an approach previously used for vegetation variables [see Beck et al., 2006], for which one would expect minima to occur during winter. Indeed, as can be seen in Fig. 3.1, virtually no samples are being filled with minima outside of winter, and predominantly in the extreme northern latitudes. While my gap-filling methodology may yield unphysical values for non-vegetation variables like SWI, I do not expect the filling of SWI to have a big influence on the final results because I do not use antecedent values of SWI. Since I do not anticipate fires during the winter, having (by necessity of gap filling) potentially unphysical values of SWI in the winter should not affect results where relevant for my analysis.

3.2.2.2 Interpolation

All datasets were interpolated to a common 0.25° spatial grid. Datasets where the original spatial resolution was higher than this were averaged; the other datasets were interpolated using nearest-neighbour interpolation to avoid smoothing local extrema [Forkel et al., 2017]. Datasets that were only available at yearly time resolution (i.e. land cover, POPD) were linearly interpolated to monthly intervals. Temporally static data (i.e. AGB) were recycled. Processing was carried out before averaging to provide monthly climatological time series, where applicable.

3.2.3 Machine learning experiments

I used RF regression (see Section 2.3) to model the relationships between BA and the driver variables (predictors). The `scikit-learn` version `0.24.1` [Pedregosa et al., 2011] RF regression implementation in `Python` was used, with hyperparameters determined using five-fold random cross-validation (CV) (see Section 2.3.1) of the eventual ALL model: `n_estimators`: 500, `max_depth`: 18, and default values for all other parameters. I found, heuristically, that limiting the number of split levels (i.e. pruning the constructed trees using `max_depth`) was necessary for the computation of SHapley Additive exPlanations (SHAP) values. This is likely due to the quadratic scaling of the computational complexity with the tree depth. The `Python` SHAP package used (version `0.37.0`) was chosen for being the official package from Lundberg et al.

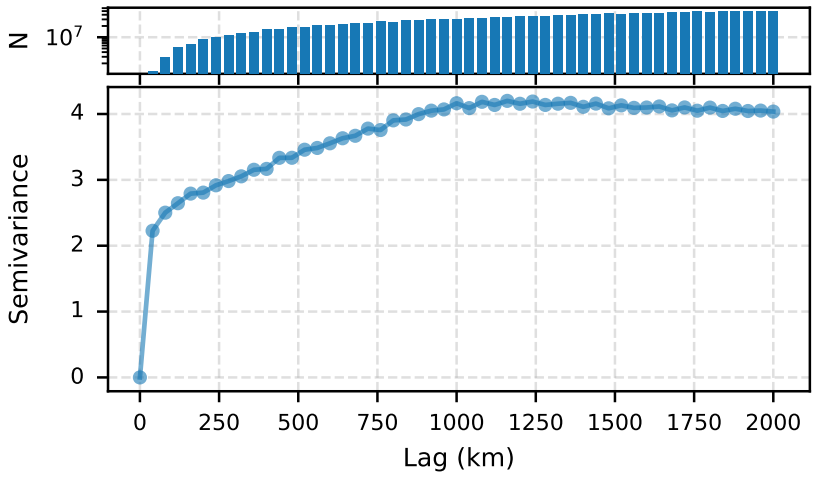


Figure 3.2: Variogram of mean GFED4 BA (June 1995 to December 2016) using all 237 373 available samples. Semivariance can be seen to increase until ∼1000 km. Note that a logarithmic scale is used for the sample counts at the top.

[2020]. The hyperparameters were only estimated once for the model containing all variables due to computational constraints.

The validation dataset was randomly sampled across space and time and comprised 30 % of the data. To estimate model generalisability, both the training dataset out-of-bag (OOB) R^2 and the validation dataset R^2 were used (see Section 2.4).

The simple CV above is only valid if there is no autocorrelation between the samples (see Section 2.4.1). I investigated the degree of spatial autocorrelation using a variogram of global GFED4 BA, which then informed a buffered leave-one-out cross-validation (B-LOO CV) procedure following Ploton et al. [2020]. This was carried out to determine how much the autocorrelation present influences the amount of potential overfitting. I did not employ the extrapolation-prevention procedure used in Ploton et al. [2020] because it led to the exclusion of significant areas like northern Australia and west Africa. This procedure is meant to prevent the examination of regions with unique conditions. Analysis of said regions necessarily implies model extrapolation since similar conditions are not found anywhere else. Thus, the applicability of the extrapolation-prevention procedure depends on the datasets considered (in the case of this analysis, see Table 3.1) and how significant the regions that would be excluded (if any) are to the analysis.

The B-LOO CV was executed with $r_{\max} = 50$ pixels (i.e. 50 grid cells), inferred from the variogram of GFED4 BA. The variogram is necessarily calculated by taking into account the distance between grid cell centres. However, given that an equirectangular long-lat projection with 0.25° grid cells throughout was used in my analysis for computational convenience, distance (in km) has to be converted to grid cells for use in the B-LOO CV procedure. At the equator, a longitudinal degree corresponds to ∼110 km, while at the poles, this distance tends towards 0 km [Budic et al., 2016]. As shown in Fig. 3.2, semivariance increases until ∼1000 km, which would correspond to ∼36 grid cells at the equator. This number of grid cells would increase moving towards the poles, which is why 50 pixels was chosen as a compromise between equatorial and high-latitude regions. Note that Budic et al. [2016] found that using a long-lat projection—with the undesirable property of unequal grid cell areas across latitudes—did not significantly influence the underlying functional relationships fitted by their statistical model compared to an equal area projection. While their study focused on the modelling of species distributions as opposed to BA, their conclusion may still be applicable.

The B-LOO CV was carried out across eight radii linearly spaced from 0 to 50 pixels. For each radius, 4000 test sample locations were chosen, and the RF model was trained on data that did not include any samples within the given radius. For each chosen test location, the test set was composed of the 12 monthly samples at that location. Due to computational constraints, the B-LOO CV was only carried out for a single set of variables.

I trained a number of different RF regression models to test explicit hypotheses about the importance of antecedent conditions on BA (see Table 3.2) using the same set of hyperparameters from the climatological time series. The initial experiment (ALL) was run using the basic set of 15 predictor variables related to climate, vegetation, and human influences on fire (Table 3.1) and included both current and antecedent values of the four vegetation indices and DD, giving 50 predictor variables in total. A second experiment (TOP15) used only the 15 most important predictors from the ALL model as a way of testing whether all the predictors were necessary and whether including so many predictors resulted in overfitting.

The choice of 15 predictors was heuristically based on the slope of the feature importance plots (see Fig. 3.19), where the importance change is minimal after 15 variables. While the use of the more rigorous recursive feature elimination with cross-validation (RFECV) would be possible in principle, this commonly makes use of the Gini importance owing to its ease of calculation, as it only considers data already seen during training. Unfortunately, this also means that RFECV fails to account for overfitting, as it only considers the training data when calculating feature importance [Meyer et al., 2019]. In contrast, the different approaches I jointly utilised to calculate a more robust feature importance metric are much more computationally demanding, making RFECV infeasible.

All of the remaining experiments used combinations of 15 predictor variables. The CURR experiment only used the current-month values of each predictor. Therefore, a comparison of the CURR and ALL experiments allowed the impact of including antecedent vegetation and moisture conditions to be evaluated. However, some of the vegetation predictors are highly correlated with one another, which could artificially decrease their importance. To test this, I ran four further experiments (15VEG_FAPAR, 15VEG_LAI, 15VEG_VOD, 15VEG_SIF) that included the 10 most important non-vegetation predictors from the ALL model, potentially including current and antecedent values of DD. In addition, each of these experiments contained one of the four vegetation predictors represented by both current (0M) and antecedent values (1M, 3M, 6M, and 9M). To disentangle the effects of antecedent DD and antecedent vegetation properties, I ran a second set of vegetation experiments (CURRDD_FAPAR, CURRDD_LAI, CURRDD_VOD, CURRDD_SIF) where each vegetation predictor was represented by both current (0M) and antecedent values (1M, 3M, 6M, and 9M) but only using current DD and the next nine most important non-vegetation factors from the CURR model. Finally, five-fold random CV was used to isolate the best combination of the vegetation predictors under the constraint that each of the five states (0M–9M) must be represented exactly once (using any of the four vegetation predictors), resulting in the BEST15 model.

In addition to the above climatological experiments, I also investigated monthly data for the

Table 3.2: The modelling experiments. Except for the ALL experiments, the other experiments included 15 predictor variables for comparability. Differences in the number of antecedent variables included in each experiment meant that a variable number of variables from the basic set were used in these experiments. For the TOP15, 15VEG_X(_MON), and BEST15 models, I used the most important variables from the ALL experiment up to the required number of 15. For the CURRDD_X models, I used the most important non-vegetation variables from the CURR model. Table 3.3 provides a detailed list of the variables included in each experiment.

Name	Nr.	Variables
ALL & ALL_NN	50	Basic set of current variables + current month and antecedent (1, 3, 6, 12, 18, 24M) values for Dry Days and vegetation indices (FAPAR, LAI, VOD, SIF)
TOP15	15	Top 15 predictors from the ALL model
CURR	15	Only current values of the basic set of 15 variables
15VEG_X (e.g. 15VEG_FAPAR)	15	Top 10 non-vegetation variables from the ALL experiment, plus current and antecedent (1, 3, 6, 9 months) vegetation index X (e.g. FAPAR)
CURRDD_X (e.g. CURRDD_FAPAR)	15	Current and antecedent (1, 3, 6, 9 months) versions of the vegetation index X (e.g. FAPAR), current DD, top 9 other variables from the CURR experiment
BEST15	15	Current and antecedent DD, one current, 1M, 3M, 6M, 9M vegetation index (drawn from the four potential vegetation indices) and five most important other variables from the basic set
15VEG_FAPAR_MON	15	Same as the 15VEG_FAPAR experiment with monthly data instead of climatological data and lightning data replaced by the next-most important non-vegetation variable

time period 11-2000–12-2029 (230 months) using the 15VEG_FAPAR_MON model. The MODIS MCD64CMQ [Giglio et al., 2018] BA dataset was used instead of the GFED4 dataset to enable an analysis of this extended period (see Table 3.1). Otherwise, this experiment uses the same variables as the 15VEG_FAPAR experiment with the exception of lightning, which was replaced with the similarly significant variable AGB (see Fig. 3.19) in order to enable processing of this longer time period. In addition to five-fold random CV as for the other models, the performance and generalisability of this model were also measured using temporal CV. Here, the model was trained on all samples excluding either all months within the years 2009–2012 (including 2012) or 2016–2019 (including 2019). Subsequently, the R^2 was measured on the years that were excluded for training. Note that unless explicitly specified, all following methodological descriptions relate to the climatological experiments as opposed to the monthly 15VEG_FAPAR_MON experiment.

Table 3.3: Variables used in the experiments. 'C' denotes current-month variables, 'all A' represents all antecedent months (1M–24M), and 1M represents one-month antecedent variables, with similar notation for other antecedent months.

	DD	SWI	MaxT	DTR	Light-ning	CROP	POPD	HERB	SHRUB	TREE	AGB	VOD	FAPAR	LAI	SIF
ALL	C & all A	C	C	C	C	C	C	C	C	C	C	C & all A	C & all A	C & all A	C & all A
ALL_NN	C & all A	C	C	C	C	C	C	C	C	C	C	C & all A	C & all A	C & all A	C & all A
CURR	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C
BEST15	C, 1M, 3M, 6M, 9M		C	C	C	C	C						C, 1M	3M	6M, 9M
TOP15	C, 1M, 3M, 9M		C			C	C					C, 1M, 3M	C, 1M	1M, 3M	C
15VEG_FAPAR	C, 1M, 3M, 6M, 9M		C	C	C	C	C						C, 1M, 3M, 6M, 9M		
15VEG_FAPAR-_MON	C, 1M, 3M, 6M, 9M		C	C		C	C				C		C, 1M, 3M, 6M, 9M		
15VEG_LAI	C, 1M, 3M, 6M, 9M		C	C	C	C	C							C, 1M, 3M, 6M, 9M	
15VEG_SIF	C, 1M, 3M, 6M, 9M		C	C	C	C	C								C, 1M, 3M, 6M, 9M
15VEG_VOD	C, 1M, 3M, 6M, 9M		C	C	C	C	C					C, 1M, 3M, 6M, 9M			
CURRDD_FAPAR	C	C	C	C	C	C		C	C	C	C		C, 1M, 3M, 6M, 9M		
CURRDD_LAI	C	C	C	C	C	C		C	C	C	C			C, 1M, 3M, 6M, 9M	
CURRDD_SIF	C	C	C	C	C	C		C	C	C	C				C, 1M, 3M, 6M, 9M
CURRDD_VOD	C	C	C	C	C	C		C	C	C	C	C, 1M, 3M, 6M, 9M			

3.2.4 Measuring predictor importance and relationships

There is no unique way to measure the importance of a given predictor on model skill in predicting BA, and it is particularly difficult to assign importance to individual predictors when there is a high degree of collinearity between them [Dormann et al., 2013, Nowack et al., 2018, Mansfield et al., 2020]. I use four techniques to infer the importance of individual predictors: Gini impurity, permutation feature importance (PFI), leave-one-column-out (LOCO), and SHAP values. The Gini importance aggregates the decrease in mean squared error (MSE) for each split involving a given predictor variable over the individual decision trees making up the RF. The PFI was calculated from five permutations of each predictor variable (validation set) using the ELI5 0.11.0 PermutationImportance (`https://eli5.readthedocs.io/en/latest/index.html`). The permutations may result in unlikely or impossible combinations of predictors, and thus the PFI approach has a known tendency to overemphasise the importance of individual variables [Hooker and Mentch, 2019]. The LOCO importance measure is estimated by repeatedly retraining the RF models, each time without one particular predictor variable. The relative importance of this predictor variable is then measured as the change in MSE on the validation dataset, where a larger drop in MSE signifies a larger significance for the variable within the dataset. The importance of correlated predictor variables may be under-emphasised in this approach since the model is retrained, and thus some of the importance associated with the removed variable may be transferred to the correlated variables during the retraining process. The SHAP value [Lundberg and Lee, 2017, Lundberg et al., 2020] is the average of marginal contributions from a series of perturbations of the predictor variables. In a similar way to the PFI, this method shares the importance amongst correlated predictor variables, which may make them appear less significant than if they were included on their own. SHAP values were computed for all validation samples. In order to create a robust composite importance metric for each predictor variable, I divided the Gini, PFI, LOCO, and SHAP metrics by the sum of their absolute values and then summed them.

Maximally significant timescales at each grid cell were calculated by weighting antecedent

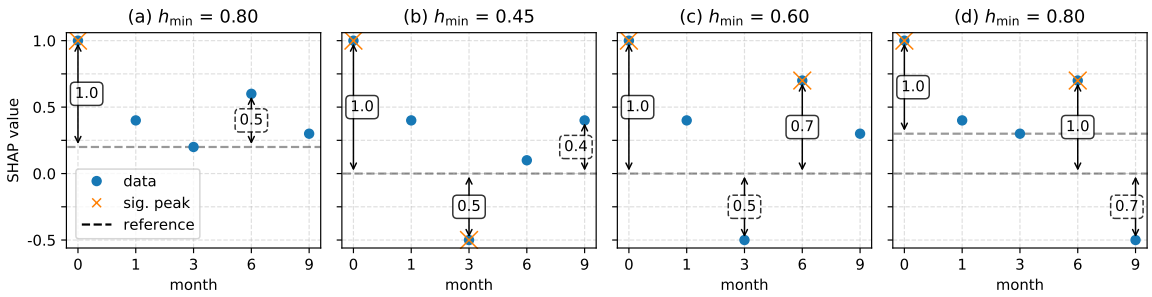


Figure 3.3: Significant peak detection algorithm. Data points are denoted by blue dots, and significant peaks are denoted by orange crosses. 'Peaks' are local extrema whose height is measured relative to the dotted reference lines, as indicated by the arrows. Their normalised height h is then calculated by dividing by the largest height. Thus, the highest peak has a height $h = 1$ by definition. Peaks are significant iff $h > h_{\min}$. Solid boxes denote significant heights, whereas insignificant heights are shown using dotted boxes. For strictly positive data (a), the reference value is the lowest adjacent local minimum. Otherwise, (see (b)–(d)), the reference value is either 0 or an adjacent extremum located between two peaks, as demonstrated by the leftmost peak in (d). Adjacent peaks are not used for height calculations to avoid spurious significant peaks, e.g. the last local minimum in (d) would be significant if its height were measured relative to the last local maximum.

months using the largest SHAP value magnitude out of all 12 months in the climatological data. The maximum SHAP value magnitudes were calculated for a predictor variable x at location ℓ on the latitude-longitude grid as follows:

$$\mathrm{SHAP}_{x,\ell} := \mathrm{SHAP}_{x,\ell,\mathrm{m}_{\max}}, \quad \text{where } \mathrm{m}_{\max} = \underset{\mathrm{m}\in\{1,2,\ldots,12\}}{\arg\max} |\mathrm{SHAP}_{x,\ell,\mathrm{m}}|. \tag{3.1}$$

These were then used to calculate the maximally significant timescales for the basis predictor variable X (e.g. FAPAR) using an average over antecedent months, a, weighted by maximum SHAP value magnitudes:

$$t_{\max,\mathrm{X},\ell} = \frac{\sum_{\mathrm{a},x} \mathrm{a}|\mathrm{SHAP}_{x,\ell}|}{\sum_{\mathrm{a},x} |\mathrm{SHAP}_{x,\ell}|} \quad \text{for } \mathrm{a}, x \in \{(0, \mathrm{X}\ 0\mathrm{M}), (1, \mathrm{X}\ 1\mathrm{M}), (3, \mathrm{X}\ 3\mathrm{M}), (6, \mathrm{X}\ 6\mathrm{M}), (9, \mathrm{X}\ 9\mathrm{M})\}. \tag{3.2}$$

Locations with too many significant antecedent months were ignored in order to visualise resulting relationships more reliably; for example, if both the current ($|\mathrm{SHAP}_{\mathrm{X}\ 0\mathrm{M},\ell}|$) and nine-month antecedent ($|\mathrm{SHAP}_{\mathrm{X}\ 9\mathrm{M},\ell}|$) magnitudes are dominant, the weighted mean month (according to Eq. (3.2)) would lie in between, which is physically meaningless. In order to mitigate this,

I designed an algorithm to detect SHAP values that differ significantly from the baseline, as illustrated in Fig. 3.3. Additionally, I also applied a range-based threshold, whereby locations were ignored if the variability of the SHAP values at location ℓ was below a threshold heuristically related to the mean BA, $\overline{\mathrm{BA}}_\ell$ (based on all BA samples):

$$\max_x(\mathrm{SHAP}_{x,\ell}) - \min_x(\mathrm{SHAP}_{x,\ell}) < 0.12 \times \overline{\mathrm{BA}}_\ell. \tag{3.3}$$

I further used accumulated local effects (ALE) plots [Apley and Zhu, 2020] to examine and interpret the coupled relationships fitted by the RF models. ALE plots are a more robust alternative to partial dependence plots (PDPs) or individual conditional expectation (ICE) plots [Molnar, 2020, Apley and Zhu, 2020]. I assessed the impact of each of the predictor variables on BA in isolation using first-order ALEs that take into account the effect of all other predictor variables. However, underlying inhomogeneities may appear when the model fits different relationships for different locations or times. I therefore tested for inhomogeneities by subsampling the dataset prior to ALE plotting to enable the visualisation of underlying relationships for a subset of locations and times. The causes of these inhomogeneities were explored using second-order ALE plots, which show the combined effect of two predictor variables on BA.

3.3 Results

In general, all models are able to predict BA using the given biophysical predictors. However, the inclusion of antecedent predictors significantly improves model performance.

3.3.1 Model performance

The ALL model, which includes all 50 variables, achieves an in-the-bag R^2 of 0.919 and OOB R^2 of 0.701 for the training dataset and an R^2 of 0.701 for the validation dataset (Fig. 3.4). Predictions on the validation set (Fig. 3.5b) show a similar geographic pattern to the observed BA (Fig. 3.5a).

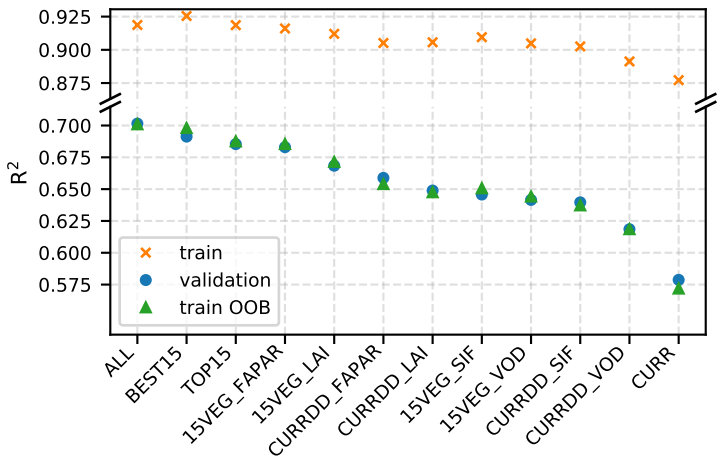


Figure 3.4: Global climatological R^2 scores for the different experiments. The CURR model is the only model that does not include antecedent conditions, and it performs much worse as a result. Despite the fact that the ALL model contains 50 predictors, while all other models contain just 15, it does not perform significantly better than the best models containing just 15 predictors (e.g. 15VEG_FAPAR). Note that although train R^2 scores are shown here, they are not indicative of model performance on unseen data, for which the shown train OOB scores should be used instead.

However, overprediction in the validation set relative to observed BA is more widespread than underprediction (Fig. 3.5c). Nonetheless, there is no bias; the ALL model predicts a mean out-of-sample BA of 2.54×10^{-3} compared to the expected 2.48×10^{-3}. The apparent overprediction is the result of plotting relative (as opposed to absolute) errors, which amplifies the fact that the ALL model does not predict very low BA accurately; out-of-sample BA predictions are no lower than 7.39×10^{-7}, while the observed BA is 0 for 85.7 % of samples. Generally, the model captures intermediate BA better than extreme BA, leading to overprediction at low and underprediction at high BA values. Thus, more samples are over-predicted because there are more values with low BA than high BA, leading to many instances of slight overprediction balanced by a few instances of comparatively large underprediction (Fig. 3.6, 3.7).

To further ensure model robustness, I also compared the PFI importances computed separately on the training and validation sets in Fig. 3.8. There is no appreciable difference between the two, which is indicative of a lack of overfitting since the model training has not unduly prioritised certain variables based on the training set [Dankers and Pfisterer, 2020]. Additionally, using the variogram shown in Fig. 3.2, I carried out the B-LOO CV as detailed in Section 3.2.3 in order to

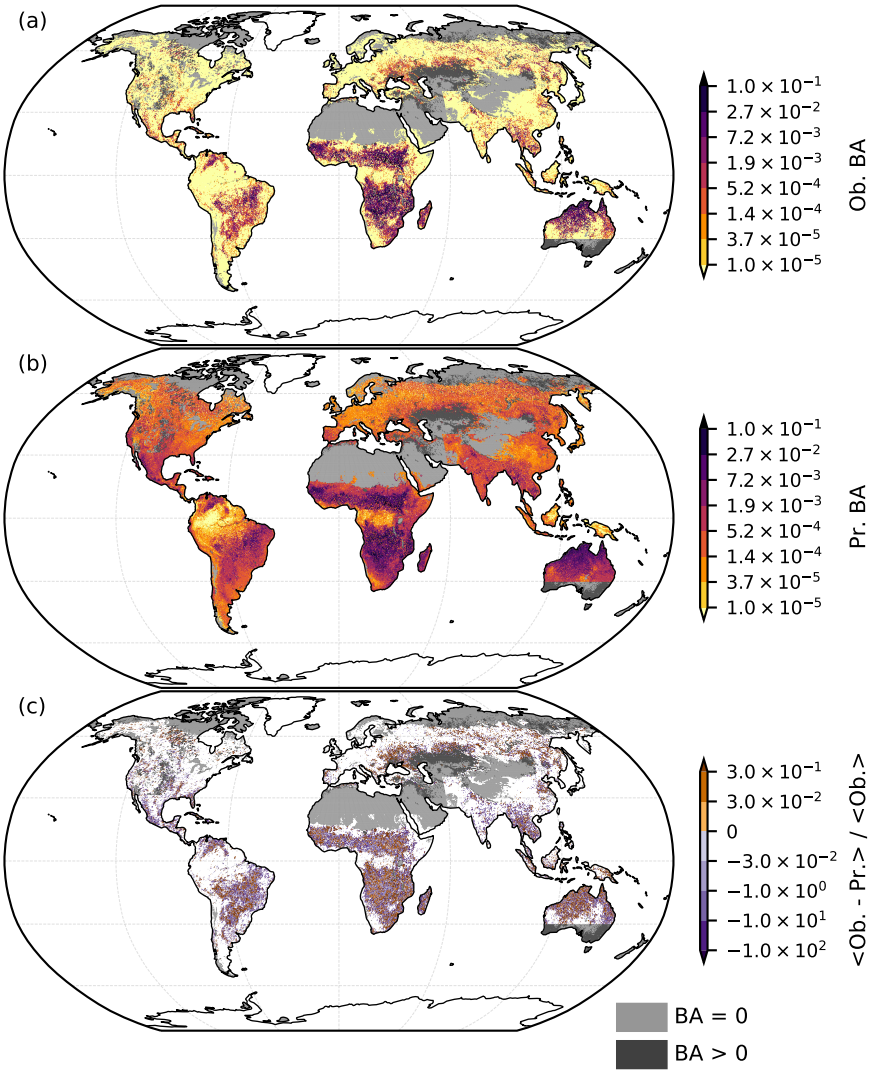


Figure 3.5: (a) Average observed (Ob.) BA fraction (1) derived from the GFED4 BA dataset [Giglio et al., 2013]. (b) Out-of-sample predictions (Pr.) by the ALL model on the validation dataset. (c) Relative prediction error of the ALL model calculated by taking the mean of the difference between observations and predictions divided by the mean observations. While the predictions in (b) are qualitatively very similar to the observations in (a), there is an overestimation of low BA. Areas with very low or 0 observed BA (a) are omitted (white) from (c) to avoid division by (nearly) 0. Despite the visual exaggeration of the errors, which are generally small, there is no overall pattern. Note that sharp data availability boundaries (e.g. in western Asia, southern Australia) are introduced by the AGB dataset. Grey shading indicates regions with fire data availability but where one or more of the other datasets are not available. Light grey indicates regions where the mean BA is 0, with dark grey representing regions with non-zero mean BA.

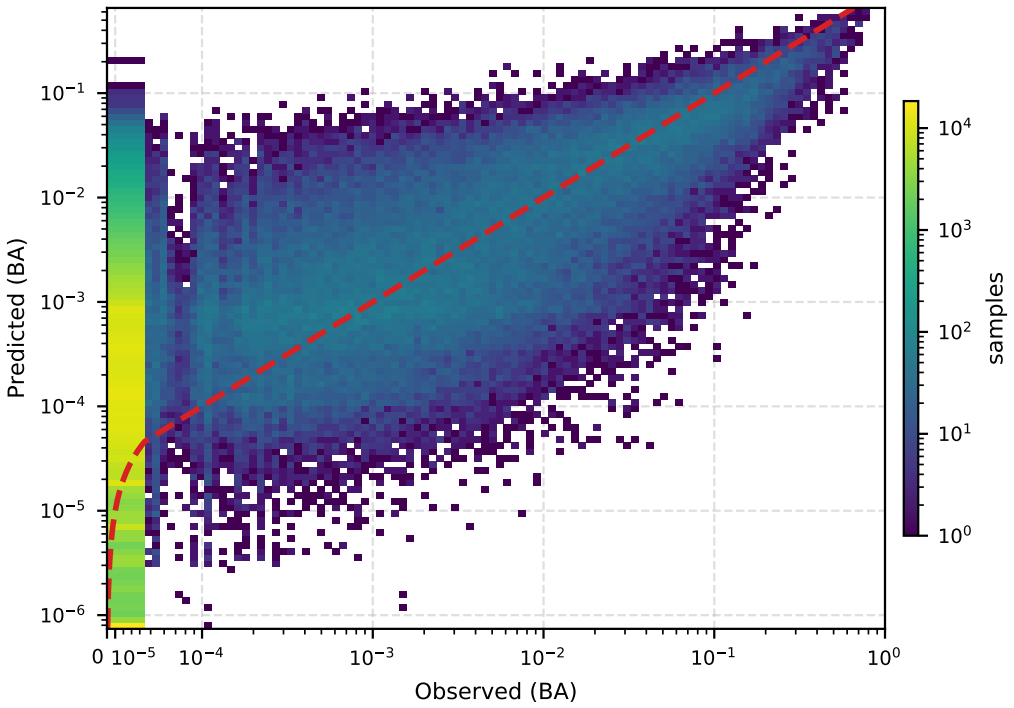


Figure 3.6: Out-of-sample BA predictions by the ALL model and corresponding observations. Note that logarithmic scales are used throughout except for the lower end of the x-axis, where a linear scale is used.

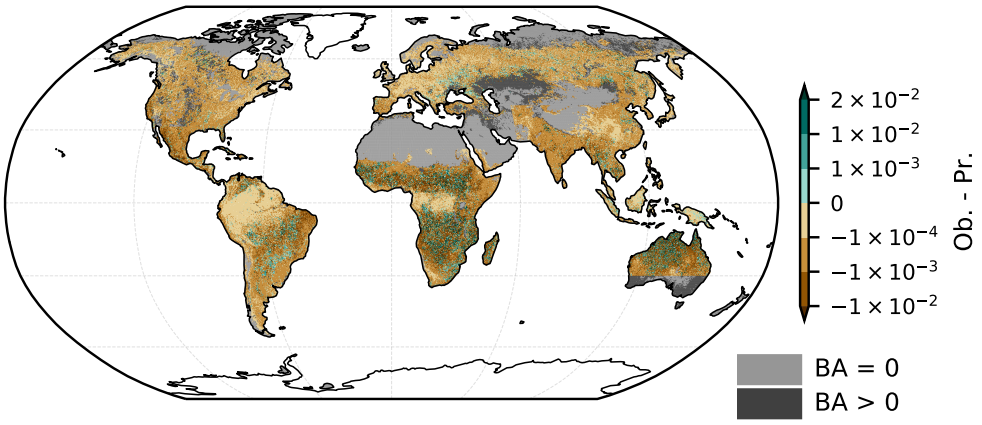


Figure 3.7: Mean difference between the out-of-sample observed (Ob.) and predicted (Pr.; by the ALL model) BA. The major spatial patterns follow the magnitude of mean BA (see Fig. 3.5a), with (on average) underprediction most prevalent in regions with large mean BA and vice versa. Note that sharp data availability boundaries (e.g. in western Asia, southern Australia) are introduced by the AGB dataset. Grey shading indicates regions with fire data availability but where one or more of the other datasets are not available. Light grey indicates regions where the mean BA is 0, with dark grey representing regions with non-zero mean BA.

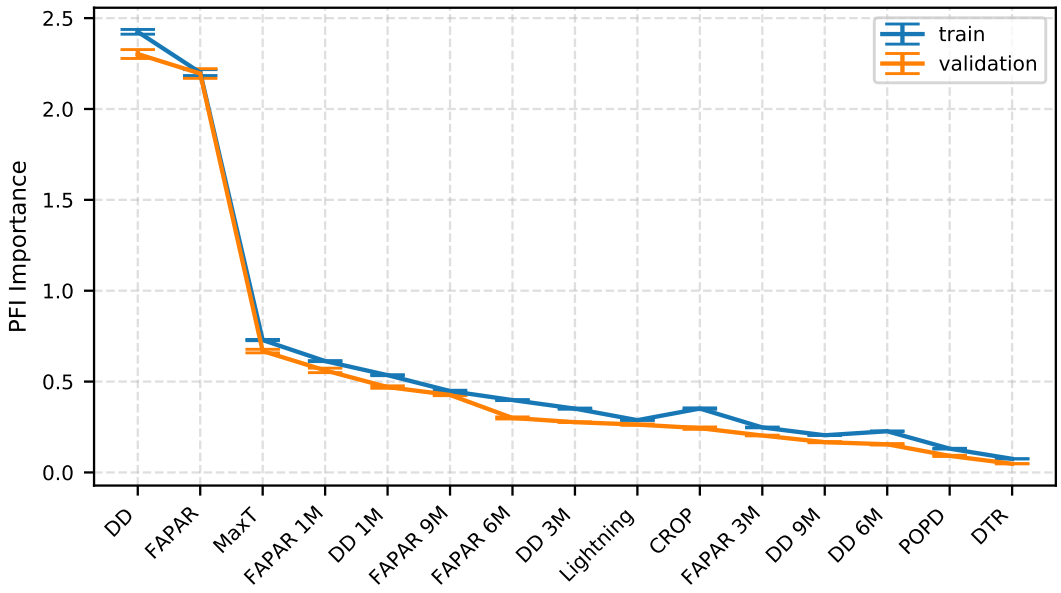


Figure 3.8: PFI importances for the 15VEG_FAPAR model computed separately on the training and validation sets. The error bars originate from the repeated shuffling of the investigated variable.

investigate the influence of spatial autocorrelation on the 15VEG_FAPAR model (see Fig. 3.9). The performance of the model drops as a larger region around the test samples is excluded (with 30 pixels corresponding to ∼900 km at the equator, which is the scale of autocorrelation identified using Fig. 3.2). Beyond ∼900 km, the R^2 score plateaus at around 0.1–0.4. Note that the extrapolation-prevention procedure in Ploton et al. [2020] was not used here because it led to the exclusion of certain key regions, such as central Africa, which typically have very large BA.

3.3.2 Importance of predictors

Climate variables and fuel-related vegetation indices have the strongest influence on BA in the ALL model (Table 3.4). Both current and antecedent conditions are important. Current DD and MaxT are ranked first and fifth, respectively, but antecedent DD also has a moderate influence (DD 1M and DD 3M are ranked 10th and 13th in importance, respectively). Similarly, although current FAPAR is the most important vegetation index (second), with both current SIF and VOD occurring in the top 15 predictors, the antecedent vegetation state also has a strong influence on BA. However, antecedent conditions >9M are not important in the ALL model. Vegetation characteristics, such as the cover of specific plant types (TREE, SHRUB, HERB) and AGB, are

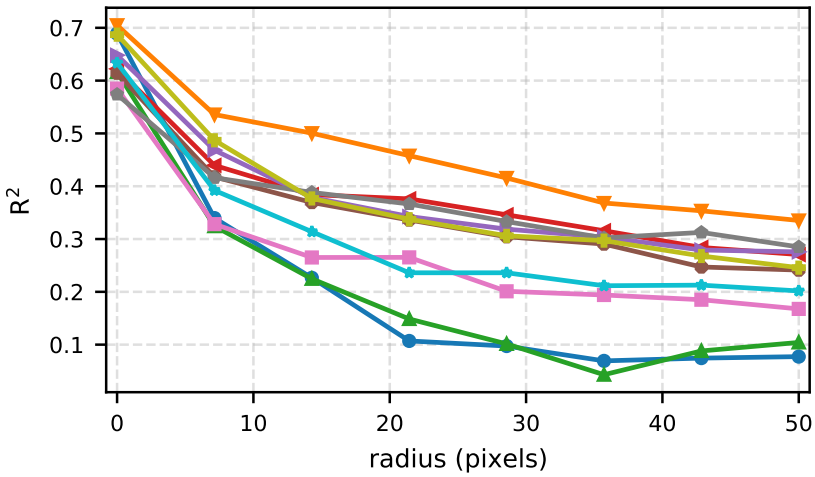


Figure 3.9: R^2 scores for BA prediction on test samples for different exclusion radii around the test samples using the 15VEG_FAPAR model. Each of the 10 lines represents the R^2 score for 400 test samples computed for the shown radii, where each individual test sample is chosen randomly and surrounded by a circular region of ignored data that is not used for training, with varying radii as shown. The disagreement between the lines is indicative of statistical uncertainty, regional variability, and potentially different degrees of model extrapolation.

only moderately important in determining BA (all ranked below the top 15 predictors). Human impacts, as represented by CROP and POPD, are also only moderately important globally for BA, ranked 8th and 15th, respectively. Natural ignitions, as represented by lightning, are only ranked 21st, suggesting that on a global climatological scale, burning is not limited by lightning.

The different importance metrics used are also in broad agreement, especially regarding the most important predictors like FAPAR and DD (see Fig. 3.10).

3.3.3 Models with fewer predictors

The model using the top 15 predictors from the ALL model (TOP15) performs only marginally worse than the ALL model, with an in-the-bag R^2 of 0.919 and an OOB R^2 of 0.688 for the training dataset and an R^2 of 0.685 for the validation dataset. This nearly equivalent performance reflects the fact that there is a high degree of correlation between several of the variables (Fig. 3.11) included in the ALL model.

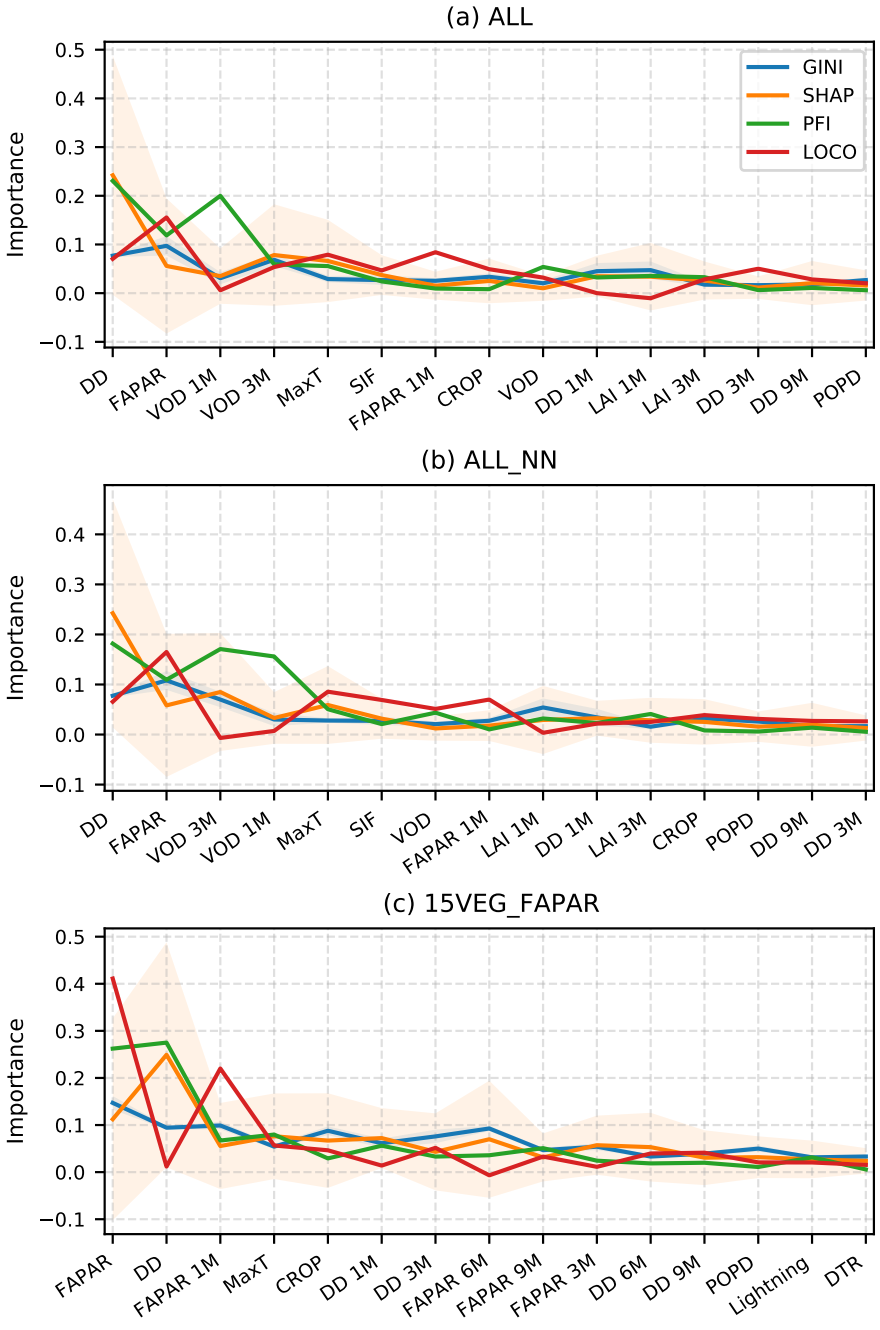


Figure 3.10: Transformed variable importance metrics (Gini, PFI, SHAP, and LOCO) for the (a) ALL, (b) ALL_NN, and (c) 15VEG_FAPAR models. The 15 most important variables (with others omitted for clarity) are sorted by their combined importance, with the most important on the left. Uncertainties using the standard deviation are indicated using shaded regions. The uncertainty magnitudes differ between the metrics due to the way they are calculated; SHAP values are calculated for every sample, Gini importances are calculated based on splits for individual decision trees, PFI calculations are repeated after permuting the original dataset, and LOCO importances are only calculated once. Therefore, based on the number of samples used for their calculation, the SHAP importances are expected to have the highest variance, followed by the Gini and then PFI importances and lastly the LOCO importances without any quantification of the error.

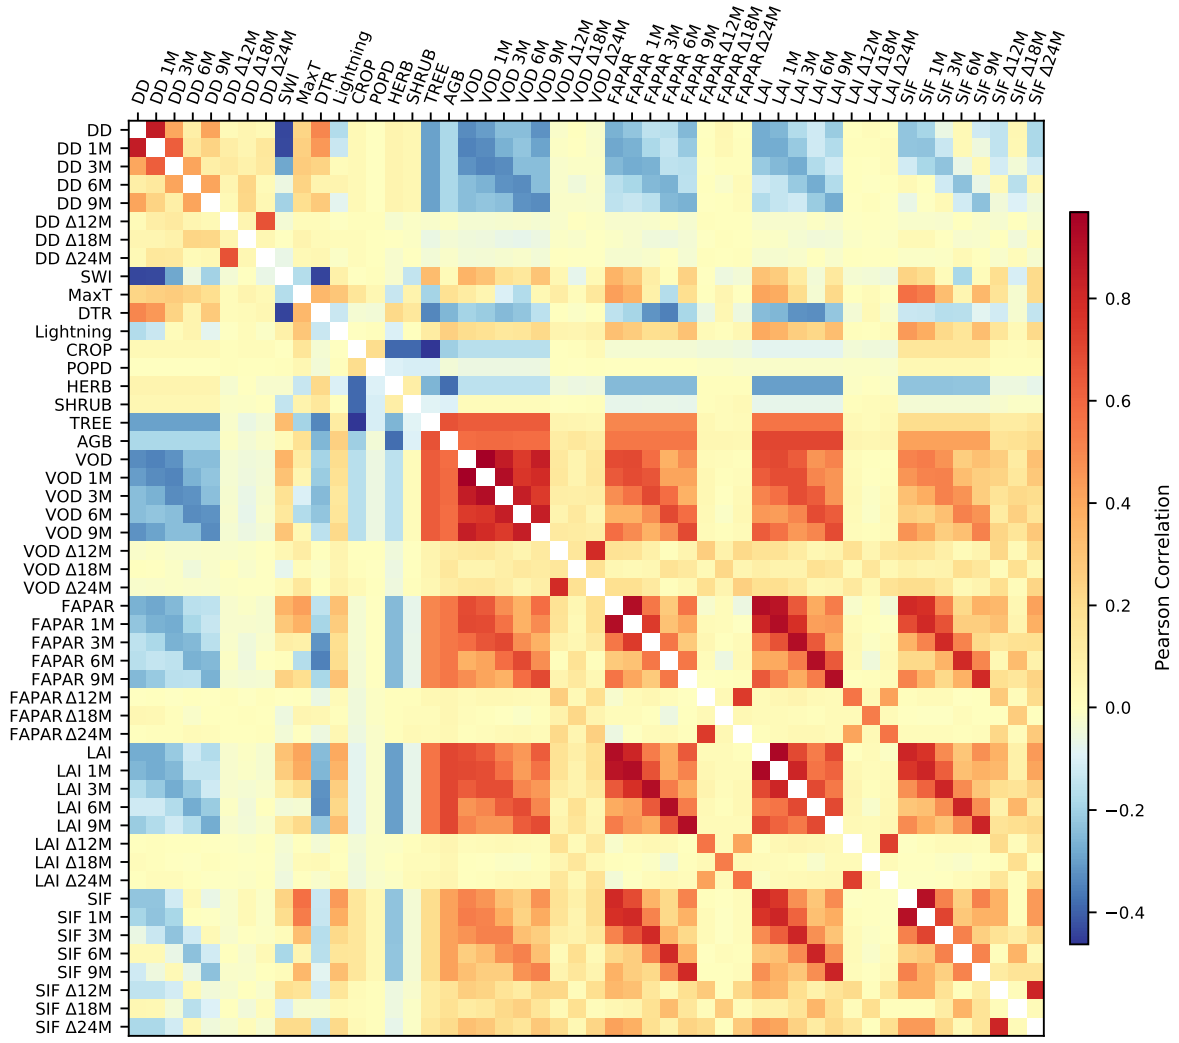


Figure 3.11: Pearson correlations between all variables used in the analysis for the time period from January 2010 to April 2015. Especially large positive correlations exist between variable pairs separated by multiples of 12 months and between FAPAR, LAI, SIF, and VOD. The largest negative correlations are found between SWI and DD (instantaneous, 12 month, 24 month), SWI and DTR, and CROP and TREE.

Table 3.4: Ranked importance of variables in the RF experiments according to the composite importance measure introduced in Section 3.2.4.

	ALL	ALL_NN	TOP15	CURR	15VEG_FAPAR	15VEG_LAI	15VEG_SIF	15VEG_VOD	CURRDD_FAPAR	CURRDD_LAI	CURRDD_SIF	CURRDD_VOD	BEST15
1	DD	DD	FAPAR	DD	FAPAR	LAI	SIF	VOD 1M	FAPAR 1M	LAI	SIF	VOD 1M	FAPAR
2	FAPAR	FAPAR	DD	MaxT	DD	DD	DD	DD	FAPAR	LAI 1M	DD	VOD	DD
3	VOD 1M	VOD 3M	MaxT	TREE	FAPAR 1M	LAI 1M	MaxT	VOD	DD	DD	MaxT	VOD 3M	FAPAR 1M
4	VOD 3M	VOD 1M	VOD 3M	VOD	MaxT	LAI 3M	CROP	VOD 3M	MaxT	LAI 3M	SIF 3M	DD	LAI 3M
5	MaxT	MaxT	SIF	SWI	CROP	MaxT	SIF 6M	MaxT	FAPAR 6M	MaxT	SIF 6M	MaxT	CROP
6	SIF	SIF	DD 9M	LAI	DD 1M	CROP	DD 1M	VOD 9M	FAPAR 3M	LAI 6M	SIF 1M	VOD 9M	MaxT
7	FAPAR 1M	VOD	LAI 1M	SIF	DD 3M	LAI 6M	DD 3M	DD 9M	CROP	HERB	SIF 9M	VOD 6M	SIF 9M
8	CROP	FAPAR 1M	VOD	FAPAR	FAPAR 6M	DD 1M	SIF 9M	CROP	FAPAR 9M	LAI 9M	TREE	AGB	DD 1M
9	VOD	LAI 1M	DD 1M	HERB	FAPAR 9M	DD 3M	SIF 3M	DD 3M	HERB	CROP	CROP	DTR	POPD
10	DD 1M	DD 1M	FAPAR 1M	DTR	FAPAR 3M	DD 9M	SIF 1M	VOD 6M	Lightning	Lightning	Lightning	Lightning	DD 9M
11	LAI 1M	LAI 3M	CROP	Lightning	DD 6M	LAI 9M	DD 6M	DD 1M	DTR	DTR	DTR	HERB	DD 6M
12	LAI 3M	CROP	LAI 3M	AGB	DD 9M	DD 6M	DD 9M	DD 6M	SWI	SHRUB	SWI	SHRUB	Lightning
13	DD 3M	POPD	VOD 1M	CROP	POPD	POPD	DTR	POPD	SHRUB	SWI	SHRUB	SWI	SIF 6M
14	DD 9M	DD 9M	DD 3M	SHRUB	Lightning	Lightning	Lightning	DTR	TREE	TREE	HERB	TREE	DD 3M
15	POPD	DD 3M	POPD	POPD	DTR	DTR	POPD	Lightning	AGB	AGB	AGB	CROP	DTR
16	SIF 9M	FAPAR 6M											
17	FAPAR 6M	SIF 9M											
18	LAI	DD 6M											
19	DD 6M	VOD 9M											
20	VOD 9M	Lightning											
21	Lightning	LAI											
22	DTR	FAPAR 9M											
23	AGB	FAPAR 3M											
24	SHRUB	DD Δ12M											
25	FAPAR 9M	DTR											
26	SIF 6M	AGB											
27	LAI 6M	LAI 6M											
28	SWI	SWI											
29	FAPAR 3M	SHRUB											
30	VOD Δ12M	SIF 6M											
31	DD Δ12M	LAI 9M											
32	SIF 3M	SIF 3M											
33	LAI 9M	VOD 6M											
34	VOD 6M	VOD Δ12M											
35	VOD Δ24M	SIF 1M											
36	VOD Δ18M	HERB											
37	SIF 1M	TREE											
38	LAI Δ24M	VOD Δ24M											
39	TREE	LAI Δ12M											
40	FAPAR Δ24M	SIF Δ24M											
41	HERB	DD Δ24M											
42	SIF Δ24M	LAI Δ24M											
43	FAPAR Δ12M	FAPAR Δ24M											
44	SIF Δ18M	SIF Δ18M											
45	SIF Δ12M	SIF Δ12M											
46	DD Δ24M	FAPAR Δ12M											
47	DD Δ18M	VOD Δ18M											
48	LAI Δ12M	LAI Δ18M											
49	FAPAR Δ18M	DD Δ18M											
50	LAI Δ18M	FAPAR Δ18M											

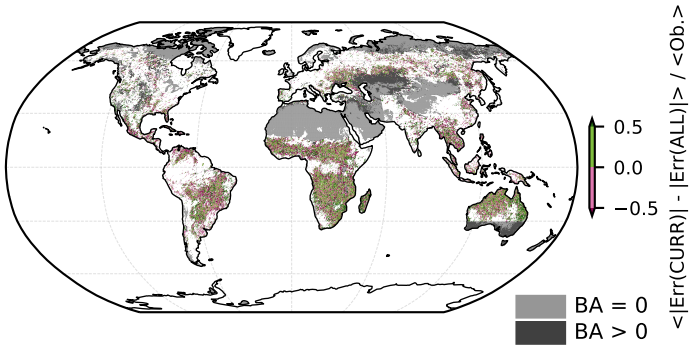


Figure 3.12: Mean change in out-of-sample prediction error between the CURR and ALL models relative to mean observations (< Ob. >). Green regions have decreased prediction error using the ALL model compared to the CURR model and vice versa for the purple regions. Areas with high BA (see Fig. 3.5a) tend to experience lower changes in relative prediction error. Areas with very low or 0 observed BA (see Fig. 3.5a) are omitted (white) to avoid division by (nearly) 0. Note that sharp data availability boundaries (e.g. in western Asia, southern Australia) are introduced by the AGB dataset. Grey shading indicates regions with fire data availability but where one or more of the other datasets are not available. Light grey indicates regions where the mean BA is 0, with dark grey representing regions with non-zero mean BA.

3.3.3.1 Importance of antecedent fuel-related predictors

The importance of antecedent fuel-related vegetation indices for predicting BA is corroborated by the results from the model that only includes predictors for the current month (CURR), where there is a large decrease in the R^2 for the validation dataset compared to either the ALL (-0.123) or TOP15 (-0.107) model. The decrease in R^2 for the training dataset is smaller (-0.042) than for the validation dataset, indicating that overfitting may be more of a problem in the CURR model than in the ALL model. Analysis of the mean out-of-sample prediction error shows that 54.7 % of grid cells are better predicted in the ALL model compared to the CURR model (Fig. 3.12). The performance improvements (from the CURR to the ALL model) also tend to have a larger magnitude than the performance decreases, contributing to the improvement in the global R^2 score. Compared to the ALL model, fuel-related vegetation properties are less important in the CURR model: VOD is the highest-ranked vegetation variable but is only fourth in importance (Table 3.4).

The four fuel-related vegetation variables included in the TOP15 model are correlated with one another (Fig. 3.11), especially on specific antecedent timescales. This suggests that it may be unnecessary to include all these variables to capture the influence of fuel build-up on BA.

Comparison of the models that only include current and antecedent conditions for one fuel-related vegetation variable (15VEG_FAPAR, 15VEG_LAI, 15VEG_VOD, 15VEG_SIF) confirms this. However, while all these models perform better than the CURR model (Fig. 3.4), only the 15VEG_FAPAR model performs similarly to the TOP15, BEST15, and ALL models. Thus, considered on its own, FAPAR was found to be the best fuel-related vegetation predictor, followed by LAI, SIF, and then VOD (Fig. 3.4).

The importance of including antecedent DD is borne out by the comparison of these four experiments and the experiments that only included current DD (CURRDD_FAPAR, CURRDD_LAI, CURRDD_VOD, CURRDD_SIF). In each case, the predictions for the same vegetation predictor variable are worse (Fig. 3.4).

The BEST15 model contains the best combination of the fuel-related vegetation predictors (current FAPAR, FAPAR 1M, LAI 3M, SIF 6M, SIF 9M), determined by optimising their timescales. This suggests that FAPAR is most important on short timescales (current, 1M), with the other vegetation properties appearing to be more useful on longer timeframes. The performance of this model, with a training R^2 of 0.925 and a validation R^2 of 0.691, is only bettered by the ALL model.

3.3.4 Current and antecedent relationships with BA

Current and antecedent states of both fuel-related vegetation properties and DD have different impacts on BA (Fig. 3.13). Current FAPAR has a negative effect on BA, while antecedent FAPAR has a positive effect on BA (Fig. 3.13a). The importance of current FAPAR changes most rapidly at intermediate levels of FAPAR. The impact of antecedent FAPAR is strongest for the preceding 1 month but persists for up to 6 months; longer lags tend to produce results more similar to the current relationship because of autocorrelation at the yearly scale. These relationships make intuitive sense: whereas high antecedent levels of FAPAR suggest that fuel availability is not a limiting factor, high FAPAR in the current month indicates that the vegetation has sufficient moisture to be actively growing and is therefore less likely to burn.

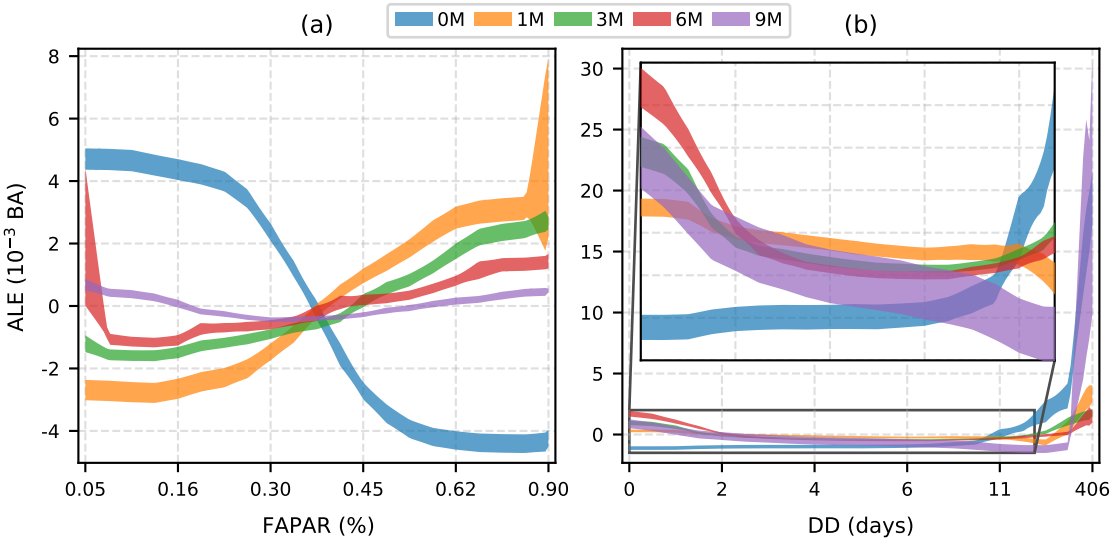


Figure 3.13: First-order ALEs for different antecedent (<1 yr) relationships with (a) FAPAR and (b) DD in the 15VEG_FAPAR model, showing the underlying relationships with BA after accounting for all other variables. The shaded regions represent the standard deviation around the mean of 100 ALEs, each using 122 567 random samples of the training data (∼10 %). Evenly spaced quantiles were used in the construction of the plots. Labels were calculated using the averaged quantiles of all the variables used. A clear difference between instantaneous and antecedent relationships can be seen in both cases, with instantaneous FAPAR limiting BA while antecedent FAPAR promotes BA and vice versa for the dry-day period. Note that the enhancement of BA due to extreme droughts (extreme dry-day periods) is apparent across time periods.

Current DD has a positive effect on BA (Fig. 3.13b), while antecedent DD has a generally negative effect, except if DD is very high when the effect becomes positive again. Whereas the positive antecedent effect of FAPAR on BA is strongest for the preceding 1M relationship and then gets weaker, the negative impact of antecedent DD gradually becomes stronger up to 9M. In contrast to FAPAR, dry conditions in the current month promote fire, whereas dry conditions in preceding months reduce vegetation growth and hence fuel build-up. However, for large DD, there was a positive relationship between DD and BA across all antecedent states on average. This positive impact of drought becomes apparent for DD > ∼10 days in the current month, whereas the threshold is higher for antecedent months: BA only increases for DD > ∼20 days for the preceding month (DD 1M) and requires > ∼40 days for DD 9M.

From Fig. 3.14a, it can be seen that the negative effects of current FAPAR are most important in wet biomes, where the limitation of fire activity due to instantaneous moisture conditions

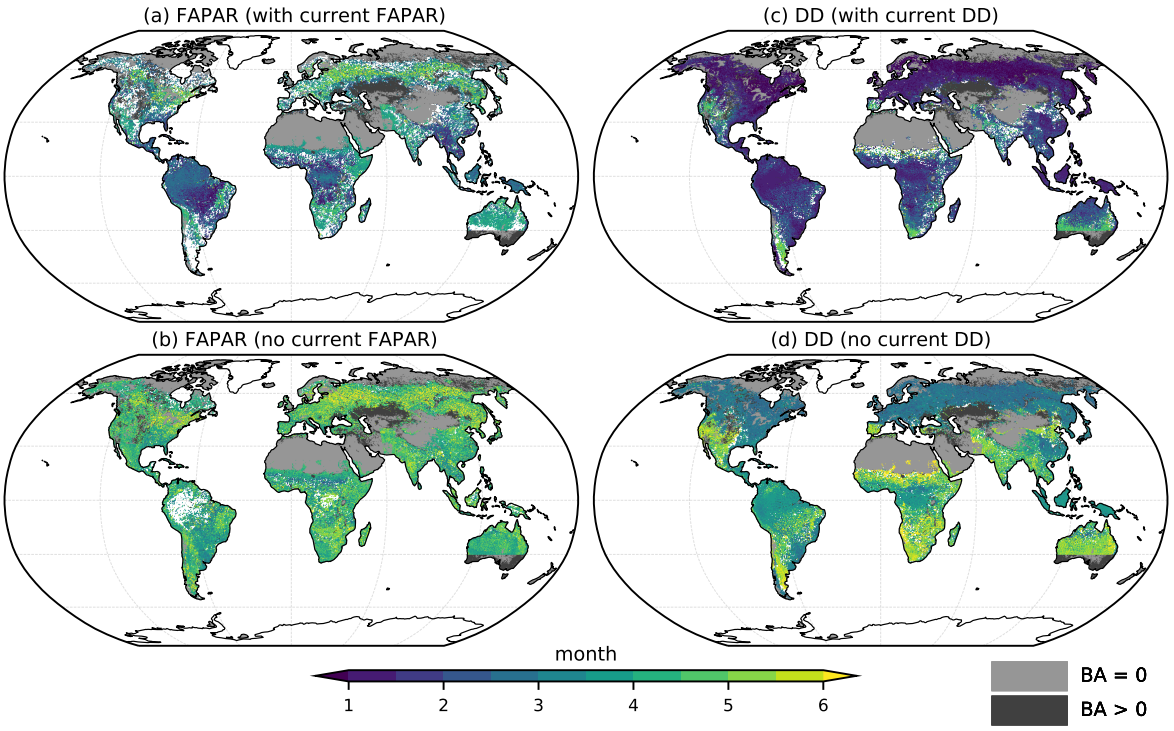


Figure 3.14: Timescales of influence of FAPAR and DD on BA. The plots show the period that is most important for determining BA from the 15VEG_FAPAR model for (a) FAPAR and (c) DD. Plots (b) and (d) show which antecedent period is most important by disregarding the influence of current conditions during plotting. Moister biomes are seen to be more influenced by current FAPAR (a), while current DD has a large influence globally (c). Note that the maps give no indication about the sign of the influence of the predictor on BA (see Fig. 3.15). Note also that sharp data availability boundaries (e.g. in western Asia, southern Australia) are introduced by the AGB dataset. Grey shading indicates regions with fire data availability but where one or more of the other datasets are not available. Light grey indicates regions where the mean BA is 0, with dark grey representing regions with non-zero mean BA.

would also be expected to be strongest. Regions where current FAPAR has a dominant influence due to the limitation imposed by high FAPAR on BA can also be seen in Fig. 3.15a. From Fig. 3.14b, it is apparent that, on average, antecedent FAPAR is most important on a ∼4-month timescale.

3.3.4.1 Consistency of relationships

Consistent relationships between current or antecedent conditions and BA are generally reproduced in all of the RF models (Fig. 3.13, 3.16, 3.17). However, the exclusion of antecedent vegetation predictors can lead to counter-intuitive relationships between the current vegetation state and BA. Although the CURR model produces the expected relationship with current

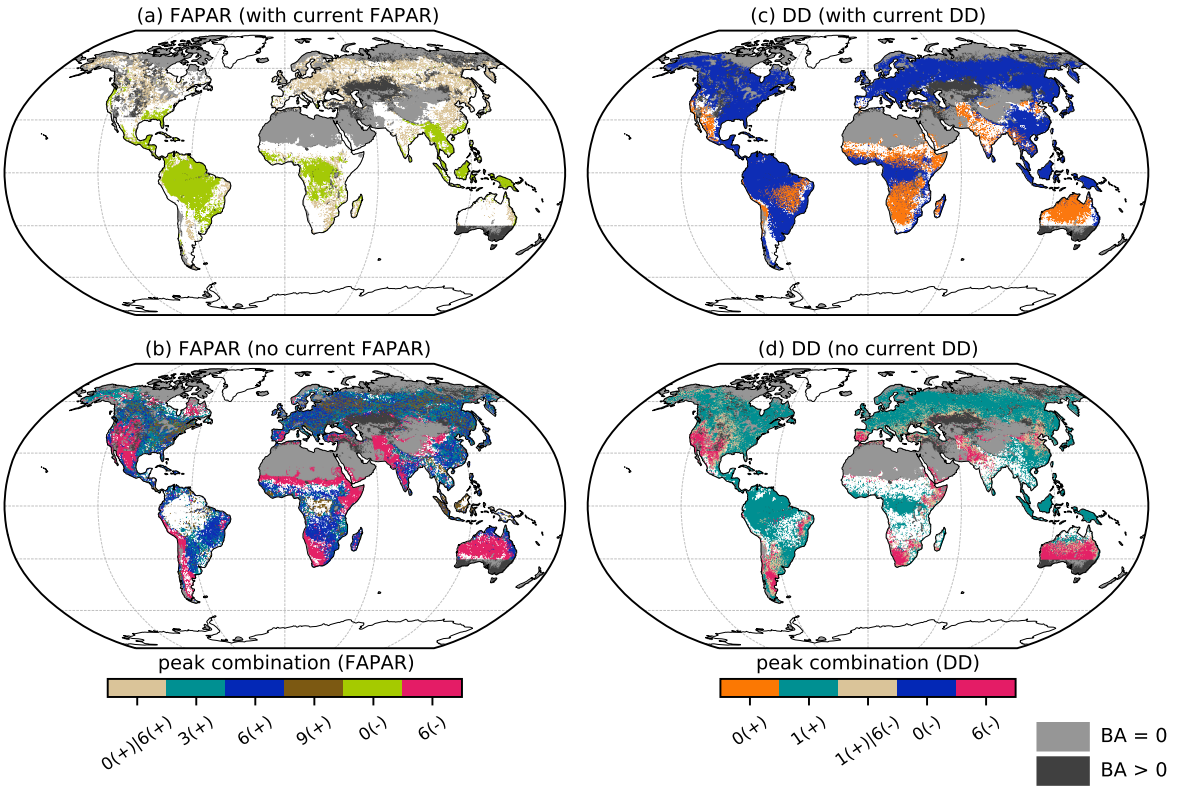


Figure 3.15: Spatial distribution of individual peak combinations for SHAP values as in Fig. 3.14. The sign of the maximum effect on BA at a certain antecedent month is indicated in parentheses after each month. The peak combinations are shown here such that their ordering has no significance (i.e. 0(+)|3(+) equals 3(+)|0(+)). Dominant antecedent periods are apparent from Fig. 3.14. Most clearly, the general limitation of BA by instantaneous DD in tropical and boreal regions is seen in (c), combined with the positive effect of instantaneous DD on burning in the remaining regions. The limitation of BA by instantaneous DD shown in (c) generally agrees with the enhancement of BA by three-month antecedent FAPAR shown in (b), as well as the enhancement by one-month antecedent DD in (d). Note that sharp data availability boundaries (e.g. in western Asia, southern Australia) are introduced by the AGB dataset. Grey shading indicates regions with fire data availability but where one or more of the other datasets are not available. Light grey indicates regions where the mean BA is 0, with dark grey representing regions with non-zero mean BA.

FAPAR (and SIF; Fig. 3.16a, 3.17a), the relationship between current LAI (and VOD) and BA is initially positive and then flat (Fig. 3.16b, 3.17b); this model does not show the expected strong negative relationship between current LAI and BA that occurs when antecedent moisture and vegetation conditions are included.

Relationships between predictors and BA were also stable when considering the 15VEG_FAPAR_MON model (Fig. 3.18), which not only uses monthly instead of climatological data, but also a different BA dataset. Using random CV, a validation R^2 of 0.501 and an OOB

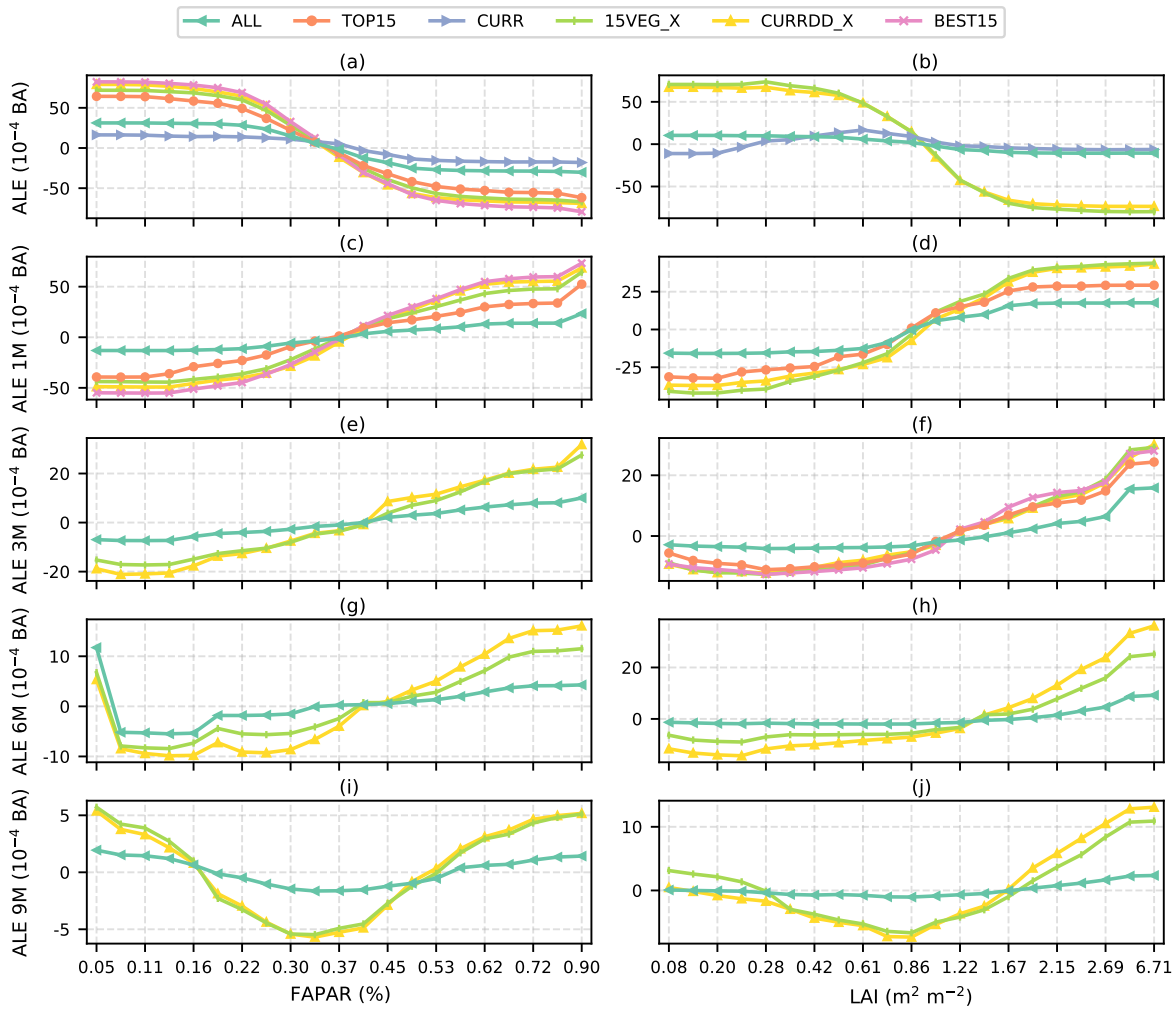


Figure 3.16: First-order ALEs for different lags (<1 yr) from all relevant modelling experiments for the relationships between BA and FAPAR (left column) and LAI (right column). Evenly spaced quantiles were used in the construction of the plots. Notably, the relationship between LAI and BA is not modelled consistently by the CURR model (b), but relationships with BA are generally consistent across models otherwise.

train R^2 of 0.498 were measured for this model. Excluding the years 2009–2012, a validation R^2 of 0.402 and an OOB train R^2 of 0.507 were measured. Excluding the final years 2016–2019, a validation R^2 of 0.435 and an OOB train R^2 of 0.505 were measured. While these R^2 scores are lower than those observed for the climatological analyses, they nonetheless demonstrate that the model is able to predict BA under multiple CV scenarios. Lower R^2 scores are also expected given the higher variance of these data. Additionally, the relationships identified by this model are highly consistent with the climatological analyses. Note also that while lightning is omitted from this experiment in contrast to the climatological 15VEG_FAPAR experiment, lightning is

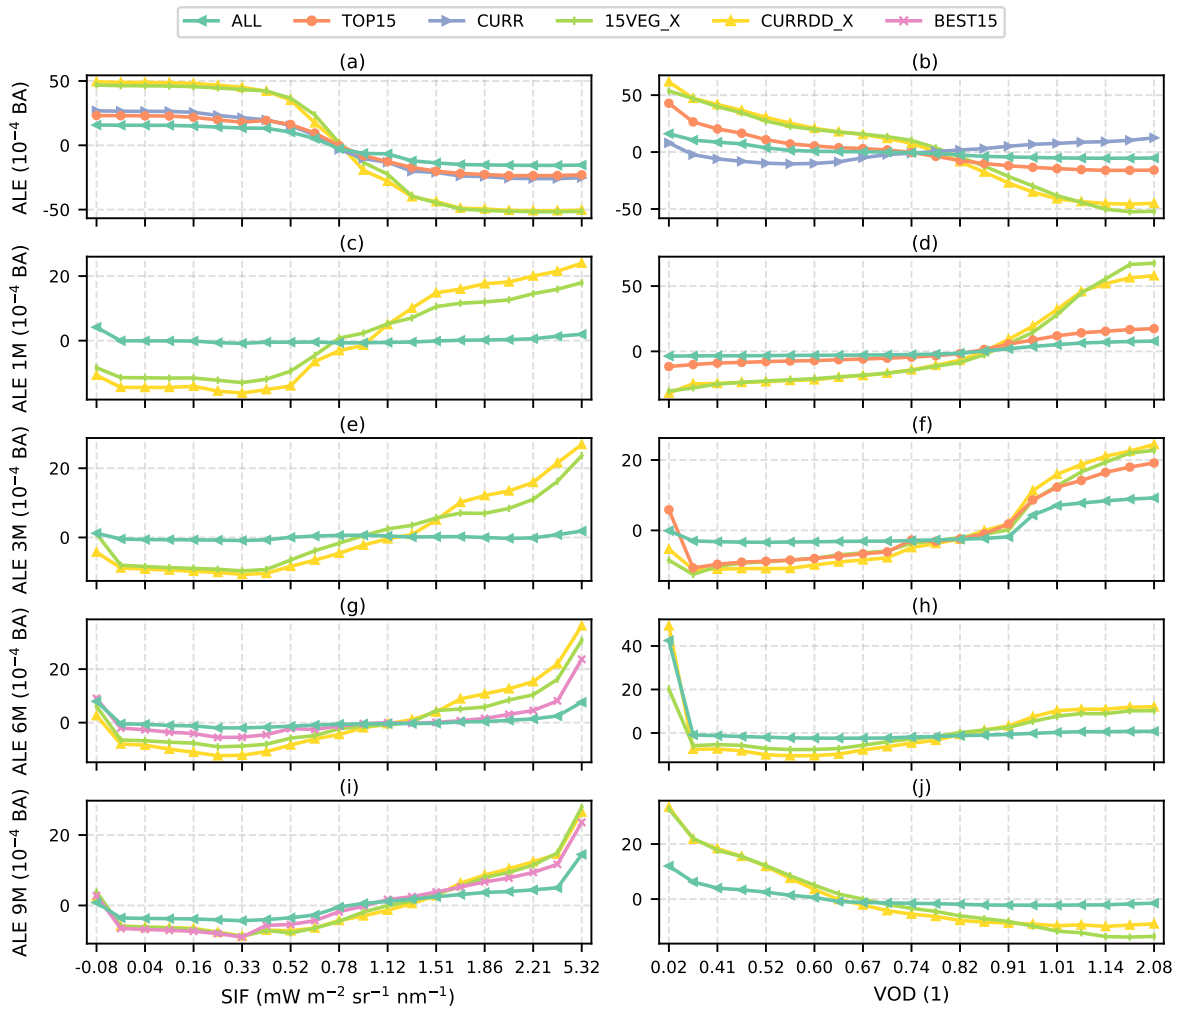


Figure 3.17: First-order LAI ALEs for different lags (<1 yr) from all relevant modelling experiments for the relationships between BA and SIF (left column) and VOD (right column). Evenly spaced quantiles were used in the construction of the plots.

also not present in the TOP15 model, which performs similarly to the ALL and BEST15 models. Furthermore, as shown in Fig. 3.19, the importance of lightning and its replacement, AGB, are very similar.

3.3.4.2 Interactions

Although it is informative to consider the impact of individual predictor variables on BA, the expression of these relationships in the real world is likely to be conditioned by interactions with other variables. For example, low values of current FAPAR are associated with high BA (Fig. 3.13a), but this association occurs only when antecedent FAPAR is high (Fig. 3.20). Low FAPAR in the current month reflects unsuitable conditions for plant growth, for example during

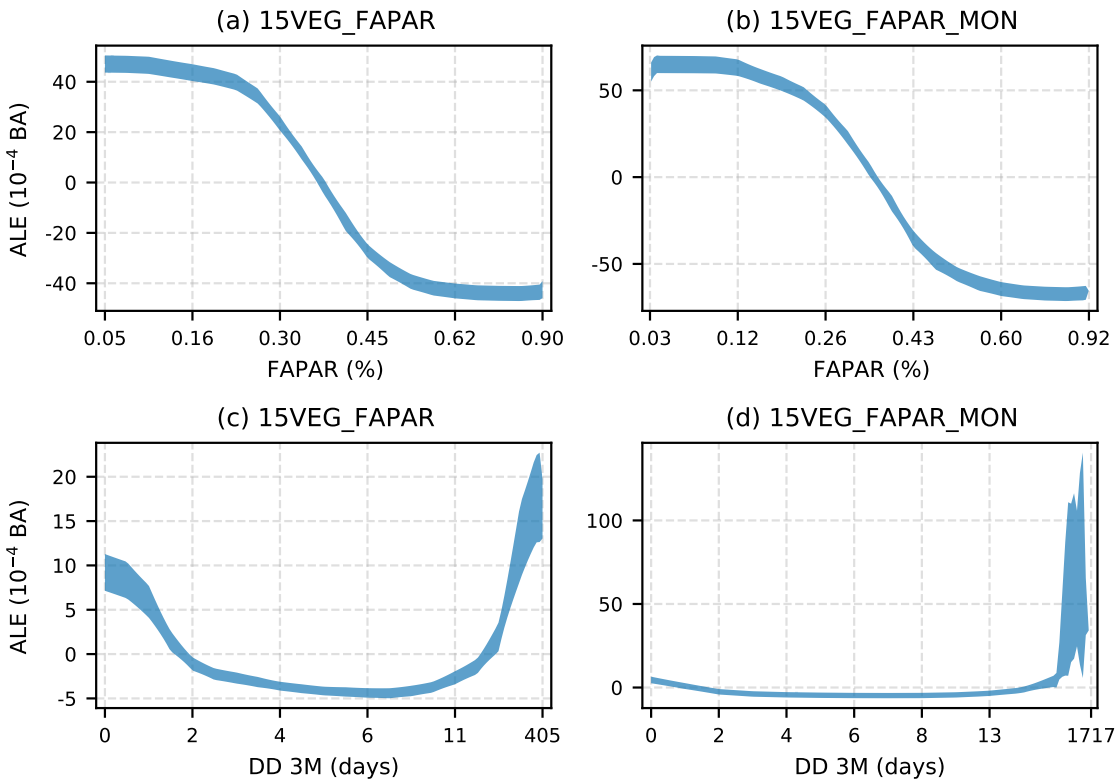


Figure 3.18: First-order ALE plots showing the effect of FAPAR (a, b) and the 3-month antecedent dry-day period (DD 3M; c, d) on BA in the 15VEG_FAPAR model (a, c) and the 15VEG_FAPAR_MON model (b, d) after accounting for all other variables. The shaded regions represent the standard deviation around the mean of 100 ALEs, each using 122 567 random samples of the training data.

the dry season. Therefore, fuel build-up during the preceding months is a prerequisite for fire to occur. The strong autocorrelation between current and preceding FAPAR values means that the occurrence of low current FAPAR coupled with high antecedent FAPAR is not widespread, being largely confined to shrublands in Africa. However, there is a significant interaction between current DD and current FAPAR (Fig. 3.21), with positive reinforcement of their mutual influence on BA when DD is high and FAPAR is low and a negative influence on BA when DD is high and FAPAR is high.

3.3.4.3 Geographically varying timescales of importance

The timescales of both fuel build-up and fuel drying are influenced by fuel type and are therefore expected to vary across biomes. Current fuel-related vegetation properties, such as FAPAR (Fig. 3.14a), have an important effect in tropical regions, particularly dry tropical regions, but

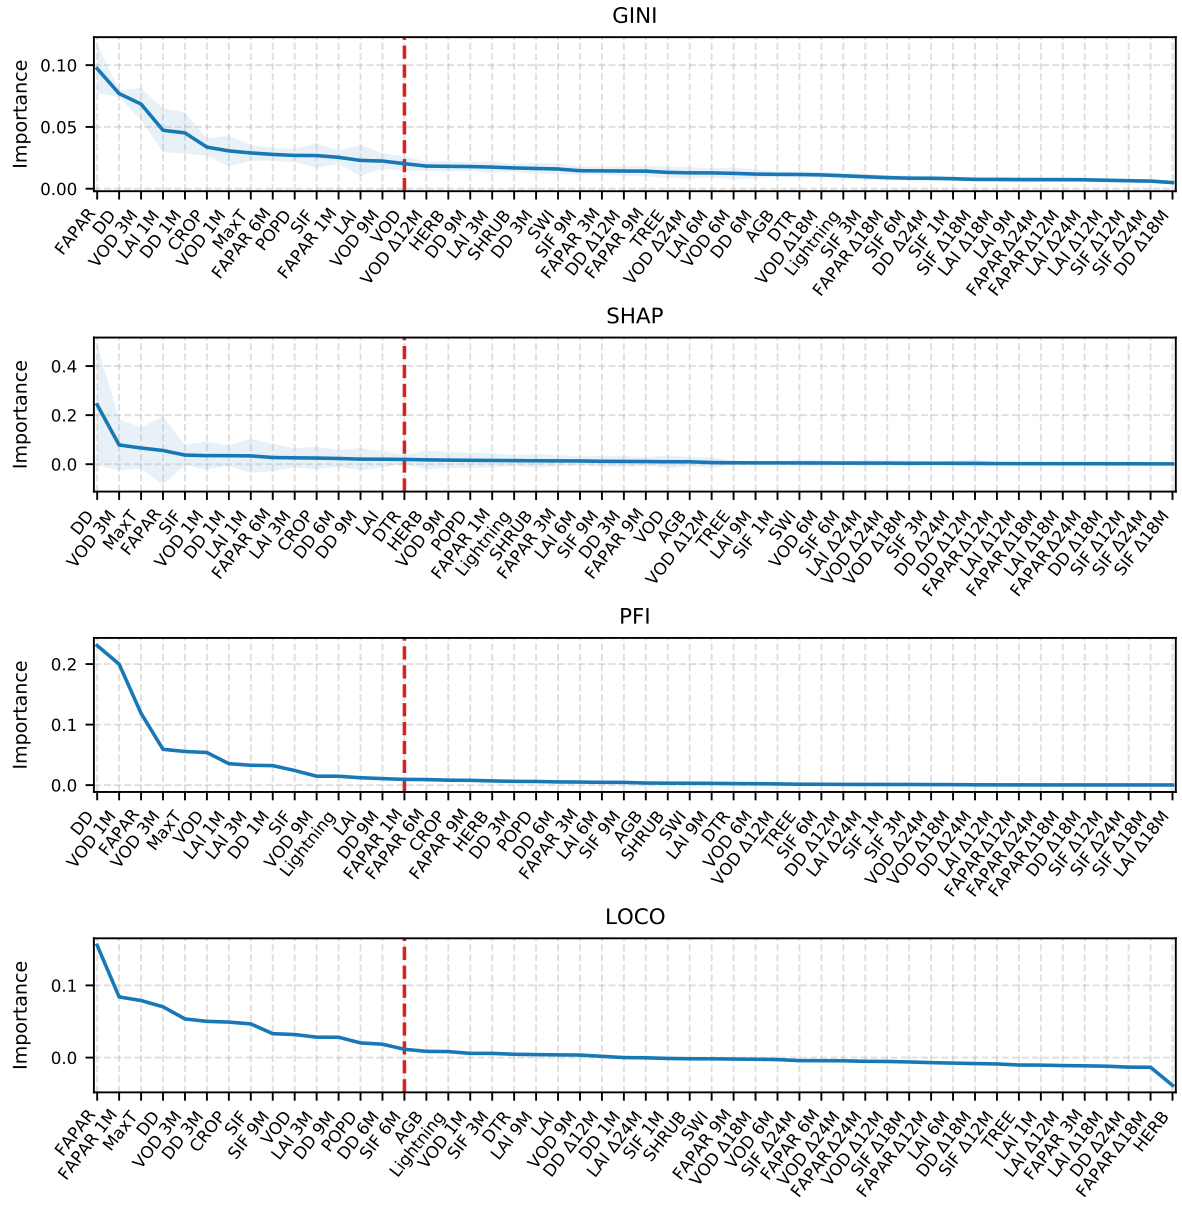


Figure 3.19: Sorted variable importance metrics (Gini, SHAP, PFI, and LOCO) for the ALL model, with the highest variable importance according to each metric on the left. The dotted red line indicates the 15th variable. Uncertainties are calculated using the standard deviation and indicated using the shaded regions.

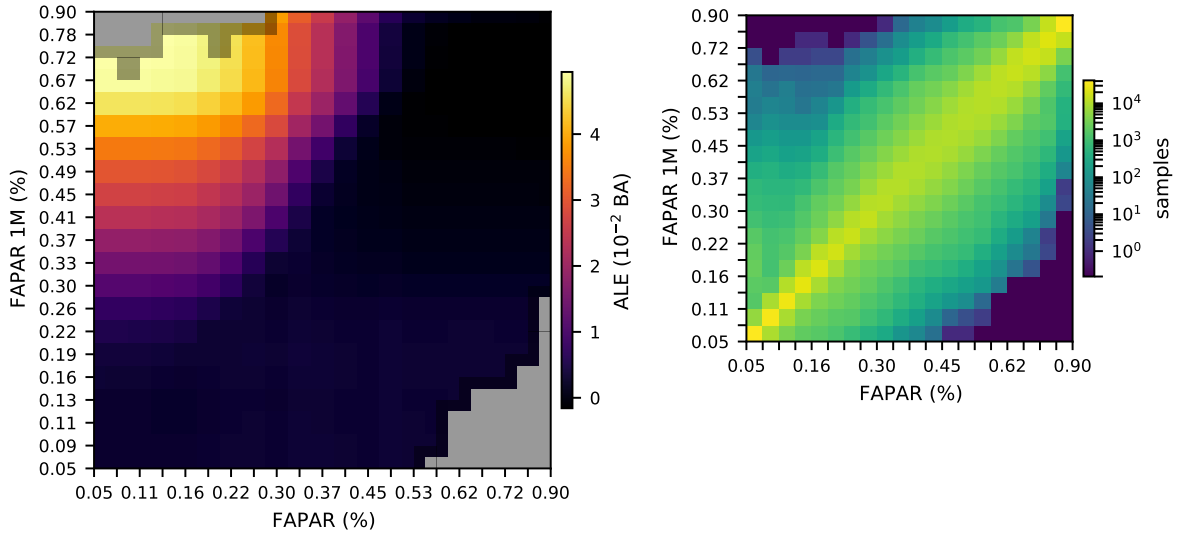


Figure 3.20: (Left) second-order ALE plot showing the combined zeroth order (mean), first order, and second order modelled effects of FAPAR and FAPAR 1M on BA from the 15VEG_FAPAR model, taking into account all other variables. Grey boxes indicate missing data. (Right) sample count matrix, which demonstrates the correlation between the variables, and thus shows that samples are unlikely to fall into the top-left or bottom-right bins. Evenly spaced quantiles were used in the construction and labelling of the plots. It can be seen that the combined effect of FAPAR and FAPAR 1M on BA is positive if FAPAR is low while FAPAR 1M is high.

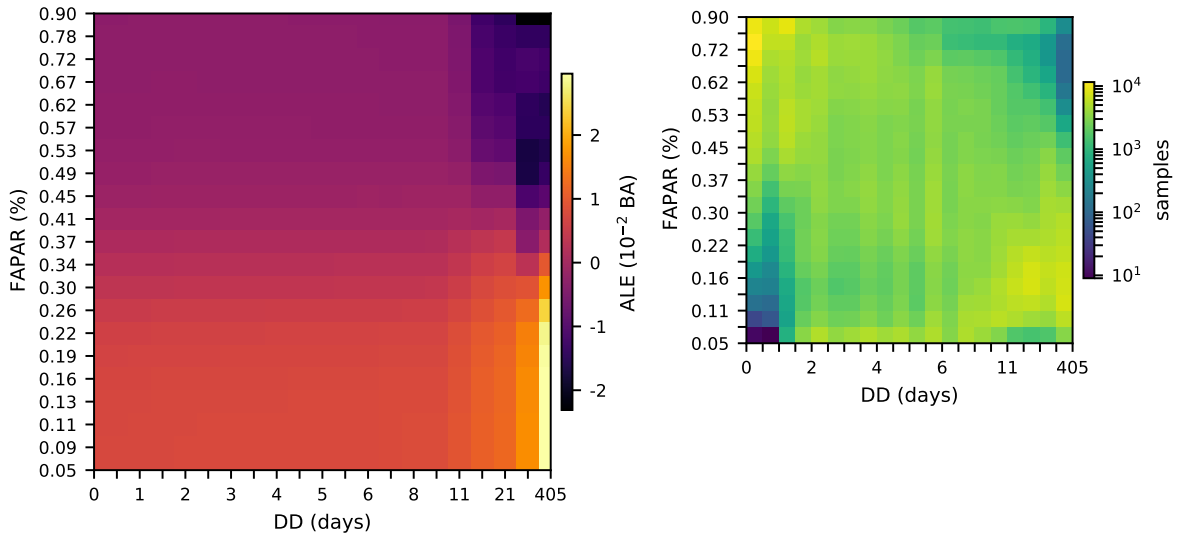


Figure 3.21: (Left) second-order ALE plot showing the combined zeroth order (mean), first order, and second order modelled effects of DD and FAPAR on BA from the 15VEG_FAPAR model, taking into account all other variables. (Right) the diagonal structure of the sample count matrix demonstrates the anticorrelation between these variables. Evenly spaced quantiles were used in the construction and labelling of the plots.

are less important in temperate forest regions. Antecedent FAPAR (Fig. 3.14b) is important in most regions, with the strongest influence from the antecedent 3–6 months. Current DD (Fig. 3.14c) is generally more important than antecedent DD, although the impact on BA varies geographically: tropical and boreal regions show decreased BA as a result of low current DD, while northwestern Australia, extra-tropical Africa, the Cerrado of Brazil, and the western USA experience increased burning as a result of high current DD (Fig. 3.15c). The timescale on which antecedent drought affects BA (Fig. 3.14d) is more variable than that for fuel-related vegetation properties, ranging from 1–3 months in boreal forests, parts of sub-Saharan Africa, and northern Australia, to ∼4 months in the tropics and ∼6 months or longer in more arid regions.

3.4 Discussion

I will first discuss the RF model prediction performance relative to earlier studies and the importance of the predictors. Next, I will consider possible limitations of the models that may affect these conclusions before finally discussing the underlying relationships with BA.

3.4.1 Model performance

Overall, the RF models were able to predict BA well using the available biophysical predictors, with the inclusion of antecedent predictors improving performance in addition to the expectedly significant climatic dry-days and temperature variables. Using a combination of regional neural networks trained on fewer variables at a coarser spatial resolution of $1° \times 1°$, Joshi and Sukumar [2021] found a global R^2 score for BA prediction of 0.36. An earlier study by Thomas et al. [2014] considered an R^2 score of 0.6 to indicate a robust prediction. My results compare favourably to both, with e.g. the ALL model achieving an R^2 of 0.701 on the validation dataset. However, it was observed that intermediate BA is captured better than extreme values. The model may struggle to predict 0 BA because the RF model consists of many smaller decision trees. All 500 individual models would have to predict 0 to yield this value overall, which does not appear to occur given the stochastic nature of the training process. Failure to capture extreme events well

is likely due to their rarity, resulting in the absence of comparable training data [see also e.g. Joshi and Sukumar, 2021].

Investigating the influence of spatial autocorrelation in a B-LOO CV analysis of the 15VEG_FAPAR model further reinforced the robustness of the RF models. Since key regions, such as central Africa, with typically very large BA, would have been excluded by employing the extrapolation-prevention procedure in Ploton et al. [2020], this was not used here. While this did enable global coverage of the B-LOO CV analysis, the RF model was forced to extrapolate to a greater extent as the exclusion radius increased, leading to overly pessimistic performance estimates. This, along with the fact that extreme events tend to be poorly captured by the model, an effect which is likely exacerbated by the dearth of training data for the largest exclusion radii in the B-LOO CV procedure, accounts for the lower end of the observed B-LOO CV R^2 values. While the RF performance did drop as more samples around the test samples were excluded from the training set, the R^2 score plateaued at around 0.1–0.4 beyond ~900 km instead of dropping to 0 as in Ploton et al. [2020], thereby indicating the robustness of my model.

3.4.2 Importance of predictors

The finding that fuel build-up on timescales longer than a year is not an important predictor of BA may initially be surprising given that fuel build-up as a result of fire suppression has been linked to large and catastrophic fires [e.g. in the USA; Marlon et al., 2012, Higuera et al., 2015, Parks et al., 2015]. The failure to detect an influence of longer-term fuel build-up on BA probably reflects the short time interval (1–2 yrs) considered for antecedent fuel build-up, far shorter than the timescales of coarse fuel build-up in these ecosystems. Since mega-fires have not been excluded from this analysis, with their BA being part of the BA datasets used, the short time interval implies that the drivers of such events cannot be disentangled satisfactorily. The seasonal differences captured by my analyses may also be unimportant in regions where long fire return intervals (FRIs) (or fire suppression) allow fuel build-up over longer periods. Wetter forests with long FRIs may be more affected by longer-term moisture deficits [Abatzoglou and

Kolden, 2013] that are not captured in the limited time period analysed. My results do not support the finding of Alvarado et al. [2020] that the total rainfall in the 24 months leading up to the end of the dry season is most correlated with BA in fuel build-up limited tropical savannahs. However, van der Werf et al. [2008] used 13 months of fuel accumulation before the peak of the fire season to investigate herbaceous fuels since this typically includes most of the precipitation contributing to the growth of annual herbaceous plant functional types (PFTs). This supports my findings regarding the importance of fuel build-up on timescales < 2 yrs. It would be worthwhile to re-examine the influence of longer timescales on BA when longer datasets are available, as, even when considering the ∼20-year-long MODIS record (which I do, using the 15VEG_FAPAR_MON model), I am strongly limited by the available data. Thus, it is not surprising that the relationships fit by the climatological models and the longer-term monthly 15VEG_FAPAR_MON model agree, thereby indicating that there is more information from the geographic patterns common to all models.

Predictability in boreal ecosystems is expected to remain very limited because the FRIs are many times longer than the currently available time series, so there is a very large stochastic component. The lower performance of my monthly 15VEG_FAPAR_MON model with over 19 years of data, in addition to the limited predictability of boreal regions found by Joshi and Sukumar [2021] despite their use of 14 years of data, both support this. The effect of El Niño/Southern Oscillation (ENSO)-driven climate variations on fire activity [Chen et al., 2017] is also likely to be dampened by the use of climatological averaging of the underlying monthly data to an average year. Furthermore, the monthly analysis that was performed (using the 15VEG_FAPAR_MON model) was also performed on variables selected using results from the climatological analyses. Additional analyses of monthly data with a more flexible set of predictor variables may therefore be able to detect sub-decadal to decadal antecedent effects.

My analyses are also impacted by the influence of previous fires on current vegetation conditions. Burnt grid cells could have a lower FAPAR, for example, as a result of prior burning within the current month. This is a problem because I am solely interested in how pre-fire

vegetation conditions affect BA. The temporal and spatial scales of the analysis are responsible for this: a monthly analysis cannot resolve processes that occur on the order of days. Further, previous fires impact the spatially averaged vegetation properties in each grid cell. On average, this impact will scale with the burnt fraction of the affected grid cell. In savannah regions of Africa and northern Australia, roughly 10 % of a given 0.25° grid cell may burn per month. This could affect the averaged values of vegetation properties used in my analyses. Analysis using a finer spatial scale would counteract this spatial smoothing by allowing burnt pixels to be ignored so that predictor values could be estimated only from unburnt pixels instead. Using a finer temporal resolution would allow the calculation of predictor variables only up to the time of burning. In practice, however, while many variables (e.g. MODIS-based vegetation variables) are available at finer resolutions, the lack of accurate and reliable fire statistics at finer scales [Abatzoglou and Kolden, 2013] limits the temporal and spatial resolution that can usefully be analysed.

Although the limitation of the spatial (and temporal) resolution of the observations could impact the realism of my models, as could the omission of variables that affect fuel build-up, the consistency of the vegetation relationships shown by all the models including antecedent conditions indicates that processes related to fuel build-up are adequately represented by the chosen set of predictors. The various importance metrics used here typically agree, especially regarding the importance of the most significant predictors like FAPAR and DD.

The TOP15 model, which used only the top 15 predictors from the ALL model, performed only marginally worse than the ALL model using all 50 predictors, indicating that the inclusion of extra predictors in the ALL model does not improve predictive capability. While this undoubtedly reflects the high correlation (and therefore little added information upon their simultaneous inclusion) between many of the variables used in the ALL model, it is also indicative of the fact that it is not necessary to include multiple fuel-related vegetation variables in order to predict BA on the global climatological scales considered here, provided that both current and antecedent conditions are taken into consideration. The removal of predictor variables is also likely to reduce

overfitting in the TOP15 model compared to the ALL model [Runge et al., 2019, Nowack et al., 2020, Joshi and Sukumar, 2021].

The 15VEG_FAPAR model performs the best out of all the models that only include a single fuel-related vegetation variable on current and antecedent timescales, indicating that FAPAR is most suitable for BA prediction on the global climatological scales investigated here. However, all four fuel-related vegetation predictors produce reasonable results, and other predictors (e.g. VOD) have been found to be important in other empirical modelling studies [e.g. Forkel et al., 2017].

The low importance of the landcover of specific plant types and AGB for BA prediction may be explained by the fact that these data are only yearly (landcover) or static (AGB). Further, it is likely that the monthly remotely sensed vegetation proxies like FAPAR, in combination with temperature and dryness, sufficiently delineate biomes for the purpose of this investigation so as to make the inclusion of land cover information redundant.

The good performance of models including FAPAR is due to the fact that the parts of the world responding most strongly to FAPAR 0M and FAPAR 1M tend to be fuel-limited, dry biomes where herbaceous vegetation is abundant. These account for the majority of global BA [Giglio et al., 2013] and therefore influence my globally averaged model performance metrics the most. Compared to VOD, the normalised difference vegetation index (NDVI) is expected to be more sensitive to variations in herbaceous as opposed to woody vegetation [Andela et al., 2013, van Marle et al., 2016]. FAPAR can be related to NDVI given certain corrections [Myneni and Williams, 1994, Andela et al., 2013, van Marle et al., 2016]. Thus, the importance of FAPAR may be due to an enhanced response to variations of herbaceous vegetation compared to the other vegetation variables used. The importance of regions experiencing frequent fires for globally averaged performance metrics is supported by previous analyses, which have found predictability in regions with infrequent fires, like boreal regions or Europe, to be poor in contrast to regions with more frequent fires [e.g. Joshi and Sukumar, 2021].

3.4.3 Underlying relationships with BA

While current FAPAR was observed to have a negative effect on BA on average, and vice versa for antecedent FAPAR, dry conditions (high DD) in the current month promote fire, with antecedent dry conditions likely to reduce vegetation growth and hence fuel build-up. These results are mostly consistent with previous studies, such as Kahiu and Hanan [2018], who found that prolonged dry-day periods limited fire activity in sub-Saharan Africa, likely due to restricted fuel availability. On average, a positive relationship between DD and BA was observed for large DD across all antecedent states. This suggests that large antecedent DD may reflect prolonged droughts extending into the current month, thereby increasing fuel flammability and promoting fire by outweighing the effect of restricted fuel availability due to lower antecedent vegetation growth.

Despite the intuitive nature of the individual underlying average relationships fit by the RF models, e.g. between DD and BA, these relationships are variable across space and time due to interactions with other variables. For example, low current FAPAR was observed to lead to high BA predominantly for high antecedent FAPAR, indicating the combined positive influence of antecedent fuel build-up and instantaneous conditions leading to low FAPAR (e.g. during the dry season). The increased BA for high DD and low FAPAR is consistent with strong drought-induced fire in low-productivity environments of sub-Saharan Africa, northern Australia, and isolated regions bordering the African tropical rainforests. These findings therefore support previous results, e.g. by van der Werf et al. [2008] for Australia, where it was found that antecedent precipitation coupled with instantaneous drying was important for fire activity. Decreased BA as a result of increased DD and increased FAPAR is likely a sign of high-productivity environments that are not fire-prone, despite occasional droughts.

There is only a moderate negative correlation between (antecedent) DD and the vegetation variables (see Fig. 3.11). Furthermore, the most important timescales vary between DD and FAPAR (see Fig. 3.14a and 3.14c), with DD being most important near the current month and FAPAR generally being more important on timescales of 3–6 months. This indicates that these

two variables (and similarly DD compared to the other vegetation variables) represent different underlying processes, with DD being more responsible for short-term fuel drying and FAPAR more responsible for fuel build-up. Compared to current-month DD, antecedent DD is less important across models (see Table 3.4). This is corroborated by the dominance of short-term influence in Fig. 3.14c. From the likely differences in underlying processes represented by the current-month FAPAR and current-month DD, it follows that antecedent DD also represents subtly different processes from the antecedent FAPAR, again corroborated by the differences between Fig. 3.14b and 3.14d.

In wet biomes, where instantaneous moisture conditions would also be expected to limit fire activity the most, the limiting effects of instantaneous FAPAR were found to be the most important (see Fig. 3.14a). This panel also has strong similarities with the results of Boer et al. [2021], with moisture-limited (dryness-limited) regions they identified corresponding broadly to the regions where instantaneous FAPAR is dominant due to the limitation imposed by high FAPAR on BA (see Fig. 3.15a). In these regions, instantaneous conditions reduce the fuel available to burn due to moisture, and antecedent conditions are less important due to the lack of a seasonal fuel build-up pattern. This then shows that in moisture-limited regions, dry events are important for enabling fire.

Fuel build-up and fuel drying timescales were found to vary across biomes, presumably due to fuel type variations. Fuel-related vegetation properties on instantaneous timescales were found to be important in tropical regions but less important in temperate forest regions. The strongest influence of antecedent FAPAR originated from the antecedent 3–6 months in most regions. Regionally, the most significant timescales of antecedent drought were found to vary more than for fuel-related vegetation properties, with 1–3 months in boreal forests, parts of sub-Saharan Africa, and northern Australia, to ∼6 months or longer in more arid regions. This supports the findings of O et al. [2020], who found that more arid regions are predominantly influenced by positive soil moisture anomalies around 5 months before fire, likely indicating increased biomass (fuel build-up), while less arid regions are more influenced by negative soil moisture anomalies

1 month before fire occurrence that indicate fuel drying. In some arid regions, e.g. the Sahel and Australia, Krawchuk and Moritz [2011] found a positive relationship between 1–2 month antecedent soil moisture and fire activity, perhaps due to rapid fuel drying or human influence (see also [Archibald et al., 2009]). This can be seen in Fig. 3.15d (peak '1(+)'), although both missing data due to multiple interfering peak combinations and other significant peaks (e.g. 6(-)) are also apparent in these regions. This positive influence of DD 1M (i.e. peak 1(+)) can be interpreted considering the influence of DD 1M on BA shown in Fig. 3.13b. There are two dryness regimes featuring a positive influence of DD 1M on BA: for low DD 1M (i.e. high one-month antecedent moisture) and for very large DD 1M. Considering that Fig. 3.13 concerns globally averaged relationships, the low DD 1M regime likely corresponds to the rapid fuel build-up and drying effect that may be present in arid regions.

3.5 Conclusions

By using random forest algorithms to model the dependence of BA on multiple climatic and biophysical variables, I have shown that antecedent vegetation conditions that influence fuel build-up and antecedent conditions that influence fuel drying significantly improved model performance when predicting BA in a given month. FAPAR was shown to be the most significant vegetation variable, and only a single vegetation variable is required for accurate BA prediction if antecedent conditions are included. Dry-day period and maximum temperature were the most significant climatic variables influencing BA. This supports previous studies that have shown that current climate and vegetation properties are important overall determinants of BA [e.g. Aldersley et al., 2011, Bistinas et al., 2014, Forkel et al., 2017, 2019a, Joshi and Sukumar, 2021, Haas et al., 2022]. The influence of antecedent climate conditions on both fuel build-up and fuel drying has also been identified as crucial in many regions [e.g. Griffin et al., 1983, Swetnam and Betancourt, 1998, Van Wilgen et al., 2000, Westerling et al., 2003, Randerson et al., 2005, Spessa et al., 2005, van der Werf et al., 2008, Archibald et al., 2009, Krawchuk and Moritz, 2011, Jenkins et al., 2020]. Indeed, the geographical patterning of BA can be linked to the spatial

variability of fuel loads and fuel moisture [Archibald et al., 2009, Boer et al., 2021].

My model-based analyses allowed the immediate and antecedent impacts of fuel loads and fuel dryness on BA to be distinguished while also allowing their relative contributions to be determined. I have further shown that current and antecedent conditions can influence BA in opposite but intuitively understandable ways: for example, wet conditions in antecedent months lead to more fuel build-up in fuel-limited regions and promote increased BA, whereas wet conditions during any given month reduce fuel dryness and thus limit BA. Furthermore, I have demonstrated that antecedent conditions $\geq$1 year are not important on a global scale. A similar conclusion was reached by Forkel et al. [2017]. Important and intuitive interactions between instantaneous and antecedent variables were captured by the RF models, for example supporting previous findings that increased antecedent productivity (FAPAR) coupled with instantaneous drying (Dry Days) promotes fire activity [e.g. van der Werf et al., 2008].

A clear contrast between fuel- and moisture-limited regions was also identified using the spatial variation of the relationship between the antecedent FAPAR and BA. The critical timescales involved varied depending on vegetation type; longer timescales ($\sim$4 months) were more important for fuel build-up in temperate regions, while recent conditions were more important for fuel drying in the tropics. The effect of vegetation variables is also biome-dependent because of differing climatic constraints. The duration of the dry-day period in the current month had the greatest impact on BA, but antecedent DD was also important, particularly in temperate regions.

Future work could re-calibrate the model hyperparameters for each set of variables to potentially improve their performance. Approaches like B-LOO CV or my combined variable importance metric could also be used more extensively to select the predictor variables and potentially tune the model hyperparameters, given that these methods are currently very computationally intensive. The findings presented here—illustrating the importance of fuel-related vegetation properties for the prediction of BA—also support calls for more detailed data about fuel, human impacts on fuel structure, and how fire-relevant plant traits are affected by

vegetation changes [Archibald et al., 2018, Forkel et al., 2019b]. Finally, the findings presented herein have the potential to improve the modelling of fire at a global scale in order to improve the way that Earth System Models depict the interactions between climate, vegetation, and fire.

Chapter 4

Development of a New INFERNO Flammability Parametrisation

Abstract

Empirical fire models represent underlying processes that control flammability using statistical relationships, in contrast to the typically less parsimonious process-based fire models that employ more detailed representations of the processes governing fire ignition and spread. Nonetheless, empirical models such as the INteractive Fire and Emission algoRithm for Natural envirOnments (INFERNO) have matched the performance of process-based fire models at the global spatial and decadal temporal scales they were designed for. However, recent analyses have shown that the phase performance of fire models is one of the areas with the most potential for improvement. With the representation of the underlying relationships between vegetation and flammability having been identified as a particular area for improvement as opposed to the climatic relationships, I developed a new flammability parametrisation for the empirical INFERNO fire model, with an emphasis on the parametrisation of fuel build-up and drying and the phase of burning. While maintaining the parsimonious nature of the standard INFERNO model, the newly developed Sigmoidal INFERNO (SINFERNO) models made use of flexible sigmoidal underlying relationships to generally improve the modelled burnt area (BA), especially in regions with typically low BA. Multiple parametrisations were investigated for dryness (dry-days and vapour pressure deficit (VPD) and precipitation) and fuel build-up (antecedent net primary productivity (NPP) and a re-parametrised carbon pool based on the Joint UK Land Environment Simulator (JULES) carbon pool). At a global scale, there were improvements to both the magnitude (normalised mean error (NME)) and the phase (mean phase difference (MPD)) of the predicted BA, with the parametrisation using VPD and precipitation and the carbon pool performing the best.

4.1 Introduction

The INteractive Fire and Emission algoRithm for Natural envirOnments (INFERNO) is an empirical fire model currently used by the community based on the parametrisation suggested by Pechony and Shindell [2009], which relates the flammability of fuel to several biophysical variables such as vapour pressure deficit (VPD), soil moisture, and precipitation (see Section 1.5.1). I redesigned the flammability calculations from the ground up in order to maintain the original parsimonious nature of INFERNO while allowing different variables to be used.

4.1.1 Application of previous results to a new flammability parametrisation

Using the insights gained from analysis of the climatological biophysical drivers of burnt area (BA) presented in Chapter 3, INFERNO's flammability parametrisation was improved as set out below, to yield the Sigmoidal INFERNO (SINFERNO) model. The previous parametrisation was enhanced by considering plant functional type (PFT)-specific relationships to variables such as (potentially antecedent) vegetation productivity and temperature. This was done to improve the magnitude of modelled BA, especially in regions with low BA, and particularly the timing of burning. As set out in Knorr et al. [2014], sigmoidal relationships to biophysical variables will be used in SINFERNO due to their ability to represent different functional forms, e.g. step-like or exponential. Also, while Satellite Observations to predict FIre Activity (SOFIA) models [Forkel et al., 2017] similarly used modified logistic functions, my approach includes an additional parameter in order to avoid discontinuities for a more physical representation. Two dryness and two fuel build-up parametrisations will be investigated, based on results presented in Chapter 3 or utilising new variables at sub-daily temporal resolution such as the VPD. While Chapter 3 did not show the relative proportion of different vegetation types to be significant for the prediction of global BA, the underlying relationships between the other investigated vegetation and climate variables varied significantly between regions. These regions were largely delineated by the type of vegetation present. Therefore, parameters in SINFERNO will be made PFT-dependent by using the following three PFT-groups: grasses, herbs, and trees, in accordance

Table 4.1: The different SINFERNO models.

Scheme	Dryness parametrisation	Fuel build-up parametrisation
A1	Dry-days	Antecedent NPP
A2	Dry-days	Carbon pool
B1	VPD and precip ('Dry Bal')	Antecedent NPP
B2	VPD and precip ('Dry Bal')	Carbon pool

with the grouping scheme used in the original INFERNO for the PFT-specific BA. This grouping was also chosen to improve interpretability while reducing computational complexity.

4.2 Methods

Two dryness parametrisations were explored, using either the dry-day period (hereafter 'A') or a combined approach using both VPD and precipitation ('B'). Fuel build-up was parametrised using either antecedent productivity (hereafter '1') or a simple litterfall model ('2'). Alternately combining these parametrisations yielded four different models: SINFERNO-A1, SINFERNO-A2, SINFERNO-B1, and SINFERNO-B2 (see Table 4.1 for an overview). To optimise model parameters with respect to observed BA data, a composite approach was taken, using basin-hopping [Li and Scheraga, 1987, Wales and Doye, 1997, Wales and Scheraga, 1999, Wales, 2003, Olson et al., 2012] (terminated by local minimisation) along with distributed Bayesian optimisation [Bergstra et al., 2013]. Basin-hopping was chosen as the primary optimisation algorithm due to its ability to efficiently find minima in high-dimensional space and its applicability to wide subject areas [Olson et al., 2012].

4.2.1 Optimised metrics

Comparison of modelled outputs to observed BA data made use of metrics previously used in Kelley et al. [2013], notably the normalised mean error (NME) and mean phase difference (MPD). Compared to the normalised mean square error (NMSE), which provides the same information as the Nash-Sutcliffe model efficiency [Nash and Sutcliffe, 1970], the NME uses mean as opposed

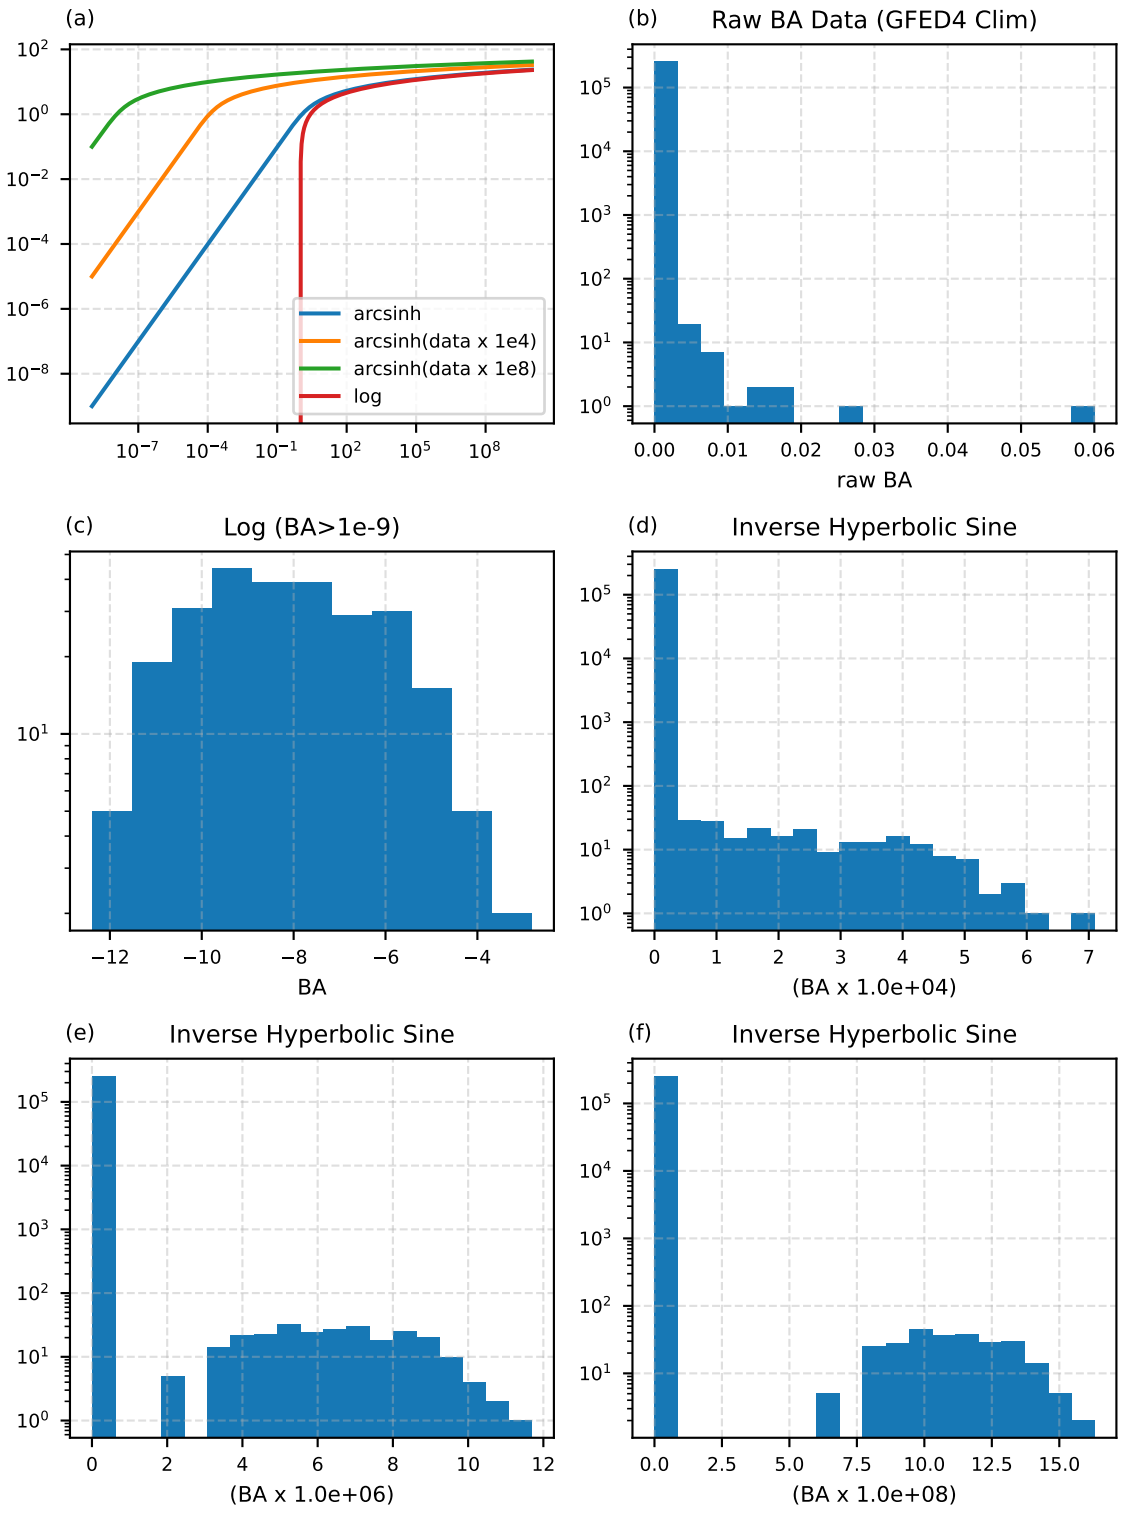


Figure 4.1: Example transformations applied to climatological GFED4 BA data. The log and arcsinh transformations are shown in (a), with their application to the raw BA data shown in (b) visualised in plots (c)–(f). The effect of varying the factor τ (see Eq. 4.1) is apparent when comparing the untransformed distribution in plot (b) to the curves shown in panels (d)–(f). With increasing τ, the discrepancy between orders of magnitude decreases as the transformed distribution qualitatively approaches the log-transformed distribution shown in (c).

to squared deviations, making the NME less sensitive to extreme values [Kelley et al., 2013]. Fourth-generation global fire emissions database (GFED4) BA data were used as the optimisation target (see also e.g. [Drüke et al., 2019]). The composite loss comprised of arcsinh-NME and MPD was minimised to determine optimal model parameters. The arcsinh-NME metric describes the NME error after data transformation using the inverse hyperbolic transformation. This was used to assign a higher relative importance to regions with BA orders of magnitude lower than the maximum observed BA (e.g. typically in boreal regions), similar to the log transform. This was done to derive a flammability parametrisation which was valid globally instead of predominantly in regions with the highest BA. As shown in Fig. 4.1, there is a clear parallel with the more commonly used log transformation depending on the factor τ used when mapping the input x to the transformed variable y, as shown in the equation below:

$$y = \operatorname{arcsinh}(\tau x). \tag{4.1}$$

For $\tau = 1$ (i.e. the standard arcsinh transformation), Fig. 4.1a shows that the log and arcsinh transforms are equivalent for inputs $x \gg 1$. However, the advantage of the arcsinh transformation as opposed to the log transformation is that it is defined for $x = 0$. By altering τ, the relative 'importance' of input values $x \ll 1$ can be greatly enhanced. During the model optimisation presented here, the value $\tau = 10^6$ was chosen heuristically based on the experiments presented in Fig. 4.1 by balancing the requirement for zeros to be represented while maintaining a shape for the remainder of the distribution similar to the logarithmic distribution (as shown in Fig. 4.1c).

4.2.1.1 Previous model implementation sources

The existing INFERNO and Joint UK Land Environment Simulator (JULES) implementations served as a starting point for some of the numerical implementations presented herein, with the underlying relationships between driving variables and flammability used by INFERNO shown in Fig. 4.2. The `Fortran` code implementing JULES (and INFERNO) can be found at the Met Office source code repository `https://code.metoffice.gov.uk/trac/jules` (registration required).

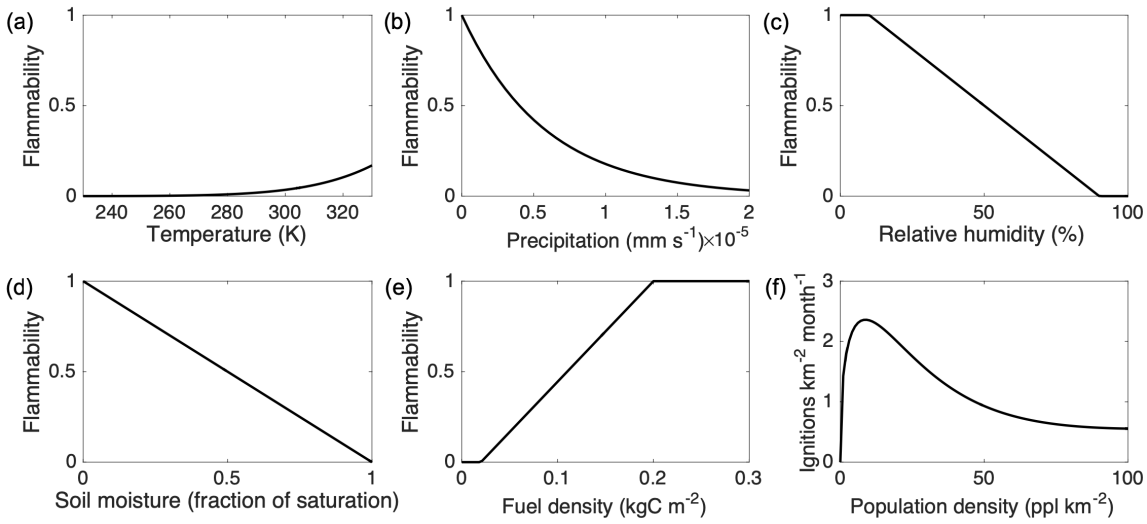


Figure 4.2: Underlying relationships between driving variables and flammability (a)–(e) and ignitions (f) used in INFERNO. Note that the relationship with population density only applies if ignition mode 3 is used. Source: Fig. A1 from Mangeon et al. [2016b], licensed under CC BY 3.0.

JULES documentation is located at `http://jules-lsm.github.io/`. Key sources regarding the current INFERNO implementation are Clark et al. [2011], Mangeon et al. [2016b], and Burton et al. [2019]. JULES version `5.4` was used, with the source code stored in `main/branches/dev/alexanderkuhnregnier/vn5.4_antecedent_fapar_5_4` at revision `r20157`, accessible via the Met Office JULES source code repository. The `u-cq574` Rose suite at revision `235535` contains the configuration used and is accessible via the Met Office JULES Rose suite repository (`https://code.metoffice.gov.uk/rosie/u/` (registration required)). While the JULES version used does include all major recent developments of the INFERNO fire module, this cannot rule out the fact that later versions may have improved upon the simulation of key vegetation properties like the fraction of absorbed photosynthetically active radiation (FAPAR), for example. However, from reviews of code changelogs (available at `http://jules-lsm.github.io/`), it appears that changes should not affect my results, with certain relevant subsequent changes (e.g. modifications to soil evaporation modelling in version `6.0`) also likely to require tuning of new parameters, which is outside the scope of this work similarly to the tuning of the numerous existing parameters within version `5.4`. Changes to the INFERNO fire module were made in version `7.0` to implement the Human Development Index (HDI)-based treatment of human impacts as proposed by Teixeira

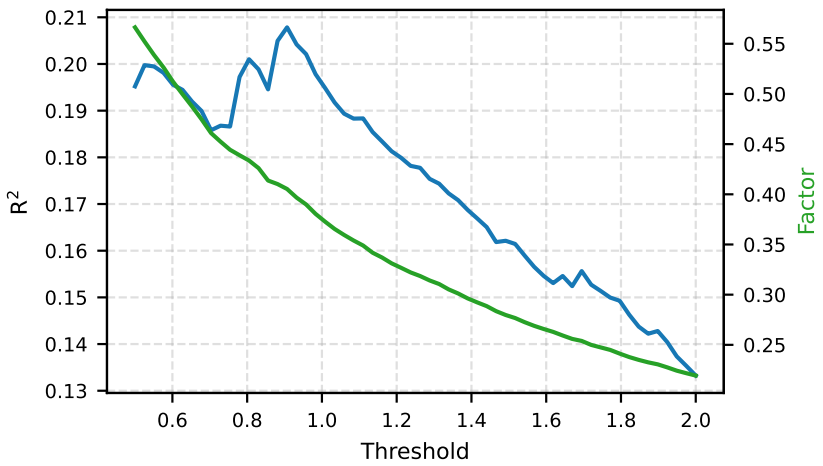


Figure 4.3: Estimation of the dry-day period calculation threshold. This threshold was used to calculate the JULES-based dry-day period. The ERA5 dry-day period, as used in Chapter 3, was used here to calculate the R^2 score (blue) and therefore the ideal scaling factor (green) for the year 2003.

et al. [2021]. However, in the `Python` SINFERNO model developed here, cropland cover was used to capture human impacts (see Section 4.2.6) instead of population density (POPD) or HDI.

4.2.2 Available JULES data

Both gross primary productivity (GPP) and net primary productivity (NPP), FAPAR, and leaf area index (LAI) are calculated within JULES for each PFT. After being spun up during 2000–2001 until the soil carbon pools converged, the JULES model was run on a 1-hour timestep for the years 2000–2016 (inclusive), with data being exported every 4 hours, as this was deemed sufficient to capture the main processes of interest without demanding too much data storage space or computational cost for analysis.

4.2.3 Dryness parametrisation

4.2.3.1 Dry-day period

The first dryness parametrisation, approach 'A', used the dry-day period calculated from 4-hourly mean precipitation data. Unlike the monthly dry-day period used in Chapter 3, which was calculated from daily mean precipitation data, this calculation yielded a dry-day period for

every 4 hours. As opposed to optimising the dry-day period threshold alongside other fire model parameters, I chose to pursue a physically accurate representation of the dry-day period using a fixed wetting precipitation threshold that is calculated in advance. This threshold was heuristically determined (see Fig. 4.3) to be ∼1.0 based on comparisons to the dry-day period derived from ECMWF Reanalysis 5th Generation (ERA5) data and calculated using the methodology from Chapter 3.

For each grid cell, JULES precipitation was first accumulated to daily sums using a moving sum before the accumulated precipitation threshold was applied:

1. Combine the JULES 'ls_rain' and 'con_rain' variables: inferno_rain = ls_rain + con_rain.
2. Calculate the daily moving sum of precipitation:

$$\text{CUM_RAIN}_t = \sum_{\max((t-\text{day_timesteps},\ 0))}^{t} \text{inferno_rain}_t, \tag{4.2}$$

   where t is the timestep index, and day_timesteps is the number of model timesteps per day.
3. Calculate the dry-day period (in timestep units, i.e. seconds) from $t = 0$ to $t = t_{\max}$ in turn:

$$\text{dry_days}_t = \begin{cases} 0 & \text{CUM_RAIN}_t > \text{threshold} \\ \begin{cases} 0 & t = 0 \\ \text{dry_days}_{t-1} + \text{timestep} & \text{otherwise} \end{cases} & \text{otherwise.} \end{cases} \tag{4.3}$$

4.2.3.2 VPD and precipitation

Dryness parametrisation 'B' used a combination of VPD and precipitation, where VPD controls fuel drying and precipitation re-wets the fuel. VPD was chosen to represent fuel drying due to its intuitive interpretation as atmospheric dryness [Fu et al., 2022] and good performance in previous fire modelling studies [Drüke et al., 2019]. The VPD was calculated following Pechony and Shindell [2009], matching the original INFERNO and Drüke et al. [2019], for example. The Dry_bal variable resulting from the combination of VPD and precipitation therefore represents

Table 4.2: Possible discrete parameter values for the 'B' dryness parametrisation (applicable to each PFT-group).

Parameter	Possible values
rain_f	0.1, 0.35, 0.6
vpd_f	50, 125, 200

a balance between the two processes of fuel drying and wetting, controlled by the PFT-group-specific factors rain_f and vpd_f, which scale the impact of precipitation and VPD, respectively. These factors represent the different behaviours of the dominant fuel classes that are associated with each PFT-group. Calculation of Dry_bal from VPD and precipitation was carried out as follows for every grid cell from $t = 0$ to $t = t_{\max}$ in turn:

$$\text{PrevDry_bal}_{i,t} = \begin{cases} 0 & t = 0 \\ \text{Dry_bal}_{i,t-1} & \text{otherwise} \end{cases} \tag{4.4}$$

$$x = \text{PrevDry_bal}_{i,t} + \text{rain_f}_i \cdot \text{CUM_RAIN}_t - (1 - e^{-\text{vpd_f}_i \text{VPD}_t}) \tag{4.5}$$

$$\text{Dry_bal}_{i,t} = \begin{cases} -1 & x < -1 \\ 1 & x > 1 \\ x & \text{otherwise,} \end{cases} \tag{4.6}$$

where CUM_RAIN_t refers to the daily moving sum of precipitation at timestep t calculated using Eq. 4.2, and i denotes the PFT-group. Note that the VPD for each PFT-group was calculated as the mean VPD over the PFTs in each PFT-group, since JULES outputs several PFT-specific variables used in the calculation of VPD.

Parametrisation 'B' required two 'discrete' sets of parameters, rain_f and vpd_f, each with three parameters, one per PFT-group. To limit computational complexity, these parameters were constrained to a preset discrete selection since for each chosen set of parameter values, the Dry_bal variable needs to be recalculated, which is orders of magnitude slower compared to BA prediction itself given a set of model parameters for several reasons:

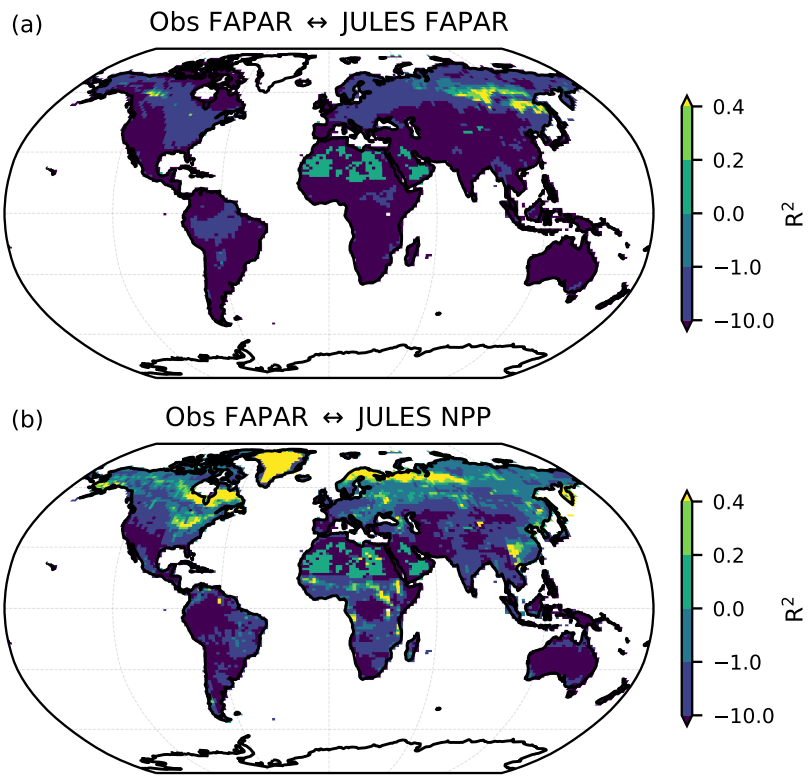


Figure 4.4: Comparison of (a) JULES FAPAR and (b) JULES NPP to observed FAPAR. Climatological variables on a monthly timescale were compared on a per-pixel basis to evaluate product similarity using the R^2.

- The Dry_bal variable is calculated using 4-hourly data instead of the weekly–monthly data used for SINFERNO. This alone accounts for over one magnitude of computational cost.
- Additional data (at 4-hourly temporal resolution) need to be loaded to enable the calculation of Dry_bal.

Optimising the rain_f and vpd_f parameters using local minimisation is therefore unfeasible. The possible values (applicable to each PFT-group separately) are shown in Table 4.2. These were heuristically derived using visualisations of the Dry_bal variable at various locations for several values of rain_f and vpd_f, which provided insight into the plausible range of each parameter. Such discrete parameters were optimised using a separate optimisation scheme (see Section 4.2.7).

4.2.4 Fuel build-up parametrisation

Fuel build-up was parametrised using either antecedent JULES NPP ('1') or a simple soil carbon pool model based on JULES litterfall ('2').

Table 4.3: Possible discrete parameter values for the '1' fuel build-up parametrisation (applicable to each PFT-group).

Parameter	Possible values
fuel_n_samples	96, 498, 900, 1302

4.2.4.1 Antecedent NPP

In Chapter 3, antecedent vegetation productivity was found to be important for build-up. Antecedent FAPAR was found to perform the best when considering a single fuel build-up proxy globally. However, as shown in Fig. 4.4, both JULES NPP and FAPAR do not correspond well to observed FAPAR. JULES FAPAR especially lacked the expected seasonality and magnitude. Therefore, I decided to use the modelled JULES NPP in the fuel build-up parametrisation (fuel build-up scheme '1'). This involved temporally shifting the NPP by a PFT-group-dependent temporal offset.

Parametrisation '1' therefore required a single discrete set of parameters, fuel_n_samples, with one parameter per PFT-group. The possible values during parameter optimisation (with the same possible values for each PFT-group) are shown in Table 4.3. These values were derived using the range of significant antecedent fuel build-up impacts observed in Chapter 3. Given the 4-hourly timestep of the data used, 96–1302 samples equate to 16–217 days, which corresponds to the range of months found to be significant in Chapter 3. Additionally, a shorter half-monthly temporal interval was investigated here, which was not possible in Chapter 3 due to the monthly data used there.

4.2.4.2 Litterfall soil carbon pool model

The second fuel build-up parametrisation (scheme '2') involved a litter pool model based on the JULES litter pool formulation as defined in Sitch et al. [2003]. The litter carbon due to leaf turnover, temperature, and soil moisture affects the soil carbon pool, together with the PFT-dependent factors litter_tc and leaf_f, which scale the rate of decomposition of the stored carbon and the fraction of leaf litter added to the pool, respectively. The aim of this re-parametrisation

Table 4.4: Possible discrete parameter values for the '2' fuel build-up parametrisation (applicable to each PFT-group).

Parameter	Possible values
litter_tc	1×10^{-10}, 5.5×10^{-10}, 1×10^{-9}
leaf_f	1×10^{-4}, 5.5×10^{-4}, 1×10^{-3}

of the existing JULES soil carbon pools is to allow fire-specific carbon dynamics to be represented. Spin-up of the soil carbon pool was performed for the year 2000 for each grid cell separately until the fractional difference between the current and previous iterations was below 0.1 %. Thereafter, the soil carbon pool was calculated for the full 17-year timespan of the analysis.

This parametrisation required two discrete sets of parameters, litter_tc and leaf_f, each with three parameters, one per PFT-group. The possible values during parameter optimisation (with the same possible values being used for each PFT-group) are shown in Table 4.4. Similar to Table 4.2, the values in Table 4.4 were heuristically derived using visualisations of the soil carbon pool variable at various locations.

4.2.5 Model structure

The flammability was parametrised using a weighted multiplication of generalised sigmoid functions. These are defined by a factor (gradient) α, a centre β, and a shape γ, with the output $s(x) \in [0, 1]$ for $\gamma > 0$:

$$s(x; \alpha, \beta, \gamma) = \left(1 + e^{\alpha\gamma(\beta - x)}\right)^{-1/\gamma}. \tag{4.7}$$

Some possible sigmoid functions centred around $\beta = 0$ are shown in Fig. 4.5. 'Weighting' of the sigmoids modulates their influence with a weight $M = 0$ yielding a constant value of $s(x) = 1$ regardless of relevant data or parameter values and $M = 1$ yielding the original sigmoid $s(x; \alpha, \beta, \gamma)$ as defined in Eq. (4.7):

$$S(x; \alpha, \beta, \gamma, M) = 1 + M \cdot [s(x; \alpha, \beta, \gamma) - 1]\,. \tag{4.8}$$

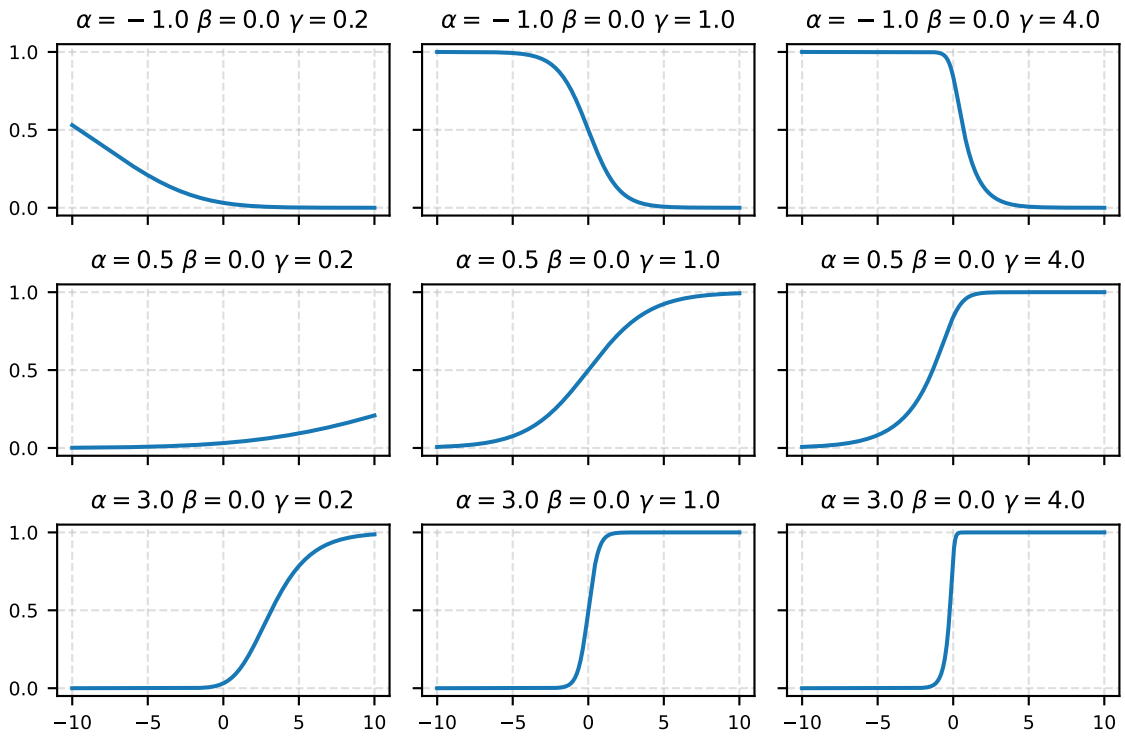


Figure 4.5: Generalised sigmoid curves (see Eq. 4.7). Different parameter values are used for the factor α and the shape γ ($\gamma > 0$). The factor parameter is seen to influence the steepness of the relationship, while the shape parameter influences the gradient asymmetry. Note that altering the β (centre) parameter amounts to a simple translation of the entire curve along the x-axis and is therefore not visualised above. The output $s(x) \in [0, 1]$, irrespective of the parameters used.

For the 13 PFTs investigated, the PFT index $\chi \in [0, 12]$. The PFT-group index ψ is defined in terms of the PFT index as follows:

$$\psi(\chi) = \begin{cases} 0, & \text{for } \chi \in [0, 4] \\ 1, & \text{for } \chi \in [5, 10] \\ 2, & \text{for } \chi \in [11, 12]. \end{cases} \tag{4.9}$$

Additional notation used to define the model structure is introduced below:

- Predicted BA (for a given grid cell): x_m, for month $m \in [0, 11]$.
- Observed BA (for a given grid cell): o_m.
- Average BA per PFT: B_χ.
- PFT fraction (for a given grid cell): $\omega_{t,\chi}$.
- Conservative averaging weights to convert model timesteps to monthly timesteps: $W_{t,m}$.

- Ignitions constant for ignition mode 1: ρ.
- Crop fraction adjustment parameter ('crop_f'): ϵ.
- Crop data at model timestep t (for a given grid cell): c_t.
- Overall scaling factor: μ.
- arcsinh-factor: τ.

The modelled flammability (for a given grid cell) at timestep t and PFT index χ is the result of the weighted multiplication of four generalised sigmoids:

$$\begin{aligned} F_{t,\chi} = & S_D(D_{t,\chi}; \alpha_{D,\chi}, \beta_{D,\chi}, \gamma_{D,\chi}, M_{D,\chi}) \cdot S_T(T_{t,\chi}; \alpha_{T,\chi}, \beta_{T,\chi}, \gamma_{T,\chi}, M_{T,\chi}) \\ & \cdot S_B(B_{t,\chi}; \alpha_{B,\chi}, \beta_{B,\chi}, \gamma_{B,\chi}, M_{B,\chi}) \cdot S_P(P_{t,\chi}; \alpha_{P,\chi}, \beta_{P,\chi}, \gamma_{P,\chi}, M_{P,\chi}), \end{aligned} \tag{4.10}$$

where $D_{t,\chi}$ represents dryness (calculated according to either scheme 'A' or 'B') at timestep t and PFT index χ, and S_D is the weighted generalised sigmoid modulating the influence of dryness. Similarly, B denotes fuel build-up calculated according to scheme '1' or '2'. Temperature is denoted by T and instantaneous NPP by P. The sigmoid parameters may only take on distinct values for each PFT-group index, with values within each PFT-group being identical by definition:

$$\alpha_{D,0} = \alpha_{D,1} = \alpha_{D,2} = \alpha_{D,3} = \alpha_{D,4} \tag{4.11}$$

$$\alpha_{D,5} = \alpha_{D,6} = \alpha_{D,7} = \alpha_{D,8} = \alpha_{D,9} = \alpha_{D,10} \tag{4.12}$$

$$\alpha_{D,11} = \alpha_{D,12}, \tag{4.13}$$

according to the PFT-groups defined in Eq. (4.9).

Predicted BA for model timestep t and PFT χ is calculated from ignitions, average PFT-specific BA, and flammability:

$$y'_{t,\chi} = \rho B_\chi F_{t,\chi}, \tag{4.14}$$

where the apostrophe indicates output on the model timestep (as opposed to monthly values). The aggregated BA over PFTs is calculated as follows using the PFT fractional coverage and

cropland data:

$$z'_t = (1 - \epsilon c_t) \cdot \left(\sum_{\chi} \omega_\chi y'_{t,\chi} \right). \tag{4.15}$$

Finally, the BA on monthly timesteps is calculated from Eq. (4.15) using conservative temporal averaging and multiplication by the overall scaling factor:

$$x_m = \mu \sum_t z'_t \cdot W_{t,m}. \tag{4.16}$$

Combining equations (4.14)–(4.16), predicted BA at monthly timesteps is calculated as:

$$x_m = \mu \sum_t (1 - \epsilon c_t) \cdot \left(\sum_{\chi} \omega_\chi \rho B_\chi F_{t,\chi} \right) \cdot W_{t,m}, \tag{4.17}$$

where the flammability is calculated according to Eq. (4.10).

Parameter optimisation was carried out using a composite metric (loss) comprised of arcsinh-NME and MPD. For a given grid cell and across 12 months, the arcsinh-NME was calculated as follows, with the denominator d in Eq. (4.18) being calculated using observations from all grid cells, however:

$$\begin{gathered} NME_a(x,o) = NME(\operatorname{arcsinh}(\tau o), \operatorname{arcsinh}(\tau x)) \\ NME_a(x,o) = \frac{\overline{|\operatorname{arcsinh}(\tau x) - \operatorname{arcsinh}(\tau o)|}}{d} = \frac{1}{d} \frac{\sum_m |\operatorname{arcsinh}(\tau x_m) - \operatorname{arcsinh}(\tau o_m)|}{12} \\ d = \overline{|\operatorname{arcsinh}(\tau o) - \overline{\operatorname{arcsinh}(\tau o)}|}, \end{gathered} \tag{4.18}$$

where $\overline{x}$ denotes the mean of x.

The MPD over 12 months was calculated as follows:

$$MPD(x,o) = \frac{1}{\pi} \arccos(\cos(P(x) - P(o))), \tag{4.19}$$

with the phase $P(x)$ being calculated as:

$$P(x) = \text{arctan2}\left(\sum_{m=0}^{11} x_m \cdot \cos(2\pi m/12), \sum_{m=0}^{11} x_m \cdot \sin(2\pi m/12)\right). \tag{4.20}$$

Denoting the observed phase at a given grid cell by $p_o = P(o)$, the overall MPD calculation for a given grid cell is as follows:

$$MPD(x, o) = \frac{1}{\pi} \arccos\left(\cos\left(\text{arctan2}\left(\sum_{m=0}^{11} x_m \cdot \cos(\pi m/6), \sum_{m=0}^{11} x_m \cdot \sin(\pi m/6)\right) - p_o\right)\right). \tag{4.21}$$

Due to the complicated nature of the composite loss compared to the squared error (i.e. model-misfit) used in other studies [e.g. Knorr et al., 2014], it was not feasible to use the analytical Jacobian during the numerical parameter optimisation as was done in [Knorr et al., 2014]. Instead, the chosen L-BFGS-B local minimisation algorithm numerically approximated the Jacobian using repeated function evaluations.

4.2.6 Cropland parametrisation

Anthropogenic influence was captured using cropland cover as a simple proxy for human fire suppression (either directly or indirectly, e.g. via agricultural land fragmentation). Although crop cover was not highly significant in the empirical analyses in Chapter 3, it was more significant than other possible indicators of anthropogenic impacts, such as population density. Therefore, I used constant ignitions (ignition mode 1), as opposed to the other modes, which take into account e.g. population density. Using the parameter $\epsilon \in [0, 1]$, the modelled BA was scaled in inverse proportion to the cropland cover in a given grid cell (using the same cropland cover dataset as in Chapter 3), according to Eq. (4.15):

$$\text{BA} \rightarrow (1 - \epsilon c_t) \cdot \text{BA}. \tag{4.22}$$

This transformation is based on the assumptions that (1) the grid cell fraction covered by agricultural land is affected by human intervention and (2) this intervention acts to decrease BA. Given these assumptions, this approach is conceptually similar to the straightforward exclusion of grid cells with high agricultural coverage used in Knorr et al. [2014].

4.2.7 Parameter optimisation

The variables described above (dry_days, Dry_bal, and soil carbon pool) were calculated using data exported from JULES every four hours. However, optimisation with respect to observed BA was carried out on a coarser timescale ranging from weekly to monthly: 42, 114, or 186 4-hour timesteps, i.e. 7, 19, or 31 days. This choice of 'averaging timescale' was optimised separately for each SINFERNO model (e.g. SINFERNO-A1). SINFERNO models were run to estimate BA after all variables were conservatively aggregated to the chosen averaging timescales. For this first-order conservative temporal aggregation, the maximum dry-day period and temperature were chosen, with the mean values used for all other variables. Coarser timescales were used because the calculation of fires on 4-hourly timescales using the empirical parametrisation presented here is unrealistic. Instead, running such an empirical model at approximately monthly timescales is "reasonable" [Pechony and Shindell, 2009], as more detailed short-term processes may otherwise be difficult to capture. Furthermore, the reference GFED4 data were only available at a monthly resolution, invalidating comparisons at much higher temporal scales.

The composite loss, MPD + arcsinh-NME with $\tau = 10^6$, was chosen to capture the agreement between modelled and observed BA. It was optimised on a monthly climatological timescale. The modelled BA was conservatively aggregated to the monthly timescale of the observed BA prior to the calculation of the composite loss. While the arcsinh transformation emphasised regions with low BA relatively more, the inclusion of the MPD metric additionally emphasised the phase of the modelled BA, irrespective of its magnitude.

Optimal parameters for each of the four possible parametrisations (see Table 4.5) were estimated using the `hyperopt` (`http://hyperopt.github.io/hyperopt/`, last accessed

Table 4.5: Number of discrete and continuous parameters being optimised across the different SINFERNO models (see Table 4.1 for an explanation of the schemes). The dry-day period parametrisation has 0 discrete parameters and thus contributes a factor of 1 to the number of possible discrete parameter combinations. The number of other possible parameter combinations can be read from Tables 4.2, 4.3, and 4.4, noting the fact that each parameter is present once for each of the three PFT-groups. The total number of discrete parameter combinations is then calculated by multiplying the number of possible averaging timescales by the number of possible discrete values for the dryness parametrisation and the number of possible discrete values for the fuel build-up parametrisation. The continuous parameters include the three sigmoid parameters α, β, and γ. There are four sigmoid curves (one for each input predictor variable), and each sigmoid parameter has three values, one per PFT-group. Thus, there are $4 \times 3 \times 3 = 36$ sigmoid parameters in total. In addition, there are three weight parameters (one per PFT-group) per sigmoid curve for an additional $4 \times 3 = 12$ parameters. Finally, including the crop_f and overall scale parameters yields 50 continuous parameters.

Scheme	Discrete parameters	Discrete parameter combinations	Continuous parameters
A1	$1 + 0 + 3 = 4$	$3 \times 1 \times 4^3 = 192$	50
A2	$1 + 0 + (3 + 3) = 7$	$3 \times 1 \times 3^{3+3} = 2187$	50
B1	$1 + (3 + 3) + 3 = 10$	$3 \times 3^{3+3} \times 4^3 = 139\,968$	50
B2	$1 + (3 + 3) + (3 + 3) = 13$	$3 \times 3^{3+3} \times 3^{3+3} = 1\,594\,323$	50

05/04/2022) [Bergstra et al., 2013] distributed Bayesian parameter optimisation method for the discrete parameters in tandem with a combined basin-hopping and local minimisation (L-BFGS-B) algorithm that was used to optimise the continuous parameters. The distinction between discrete and continuous parameters was made for computational reasons, as the computationally demanding calculation of certain quantities, e.g. the soil carbon pool, needs to be repeated for each combination of the relevant parameters. All gridpoints, not just a subset, were considered simultaneously in this optimisation, which is possible due to offline runs including only the updated fire module described here, as opposed to e.g. more complicated models and online runs including vegetation feedbacks used in Drüke et al. [2019].

4.2.8 Model analysis

4.2.8.1 Null-model analysis

Following Kelley et al. [2013], null-model analysis was used to compare the performance of the new parametrisation to the performance that would be achieved if:

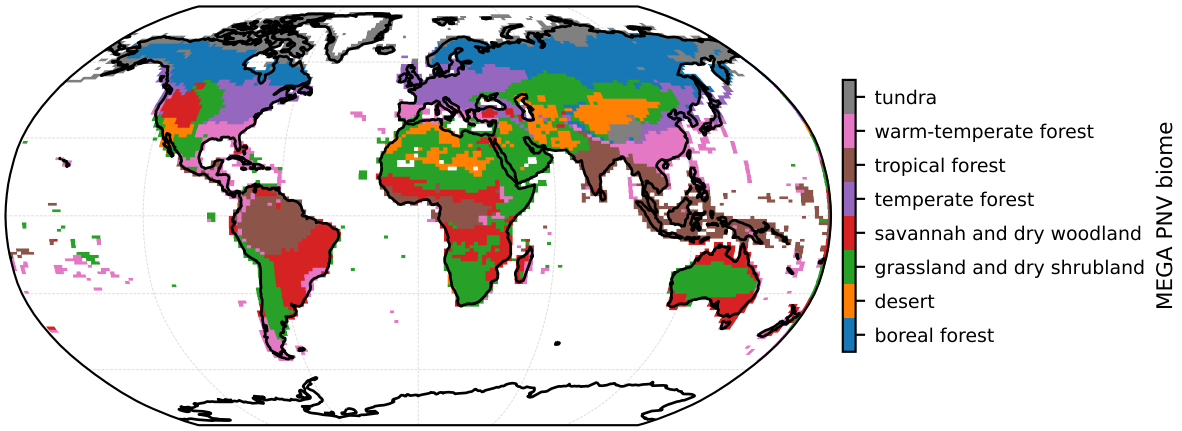


Figure 4.6: Mega PNV biomes used for regional analysis of model output (see Hengl et al. [2018], Hengl [2018]).

1. BA predictions were replaced by the global mean observed BA (this corresponds to an NME value of 1),
2. BA predictions consisted of randomly resampled observed BA.

For the random replacement, all grid cells and times were given a randomly sampled (with replacement) observed BA before calculating the NME and MPD errors. This procedure was repeated 10 000 times.

4.2.8.2 Regional analysis

In addition to the analyses and comparisons of model output on a global scale, regional analyses were made to determine the areas where the most significant changes in model performance occurred. These were carried out using regions delimited by potential natural vegetation (PNV), derived from Hengl et al. [2018] and shown in Fig. 4.6. The original biomes mapped by Hengl [2018] were converted to Mega biomes using the conversion table provided by Hengl [2018] as follows:

1. Convert the provided PNV tif file to netCDF.
2. Apply mapping from biome to Mega biome for each grid cell.
3. Regrid the resulting Mega biomes to the N96e grid using the 'mode' resampling method, which yields the most common Mega biome within each new (larger) grid cell.

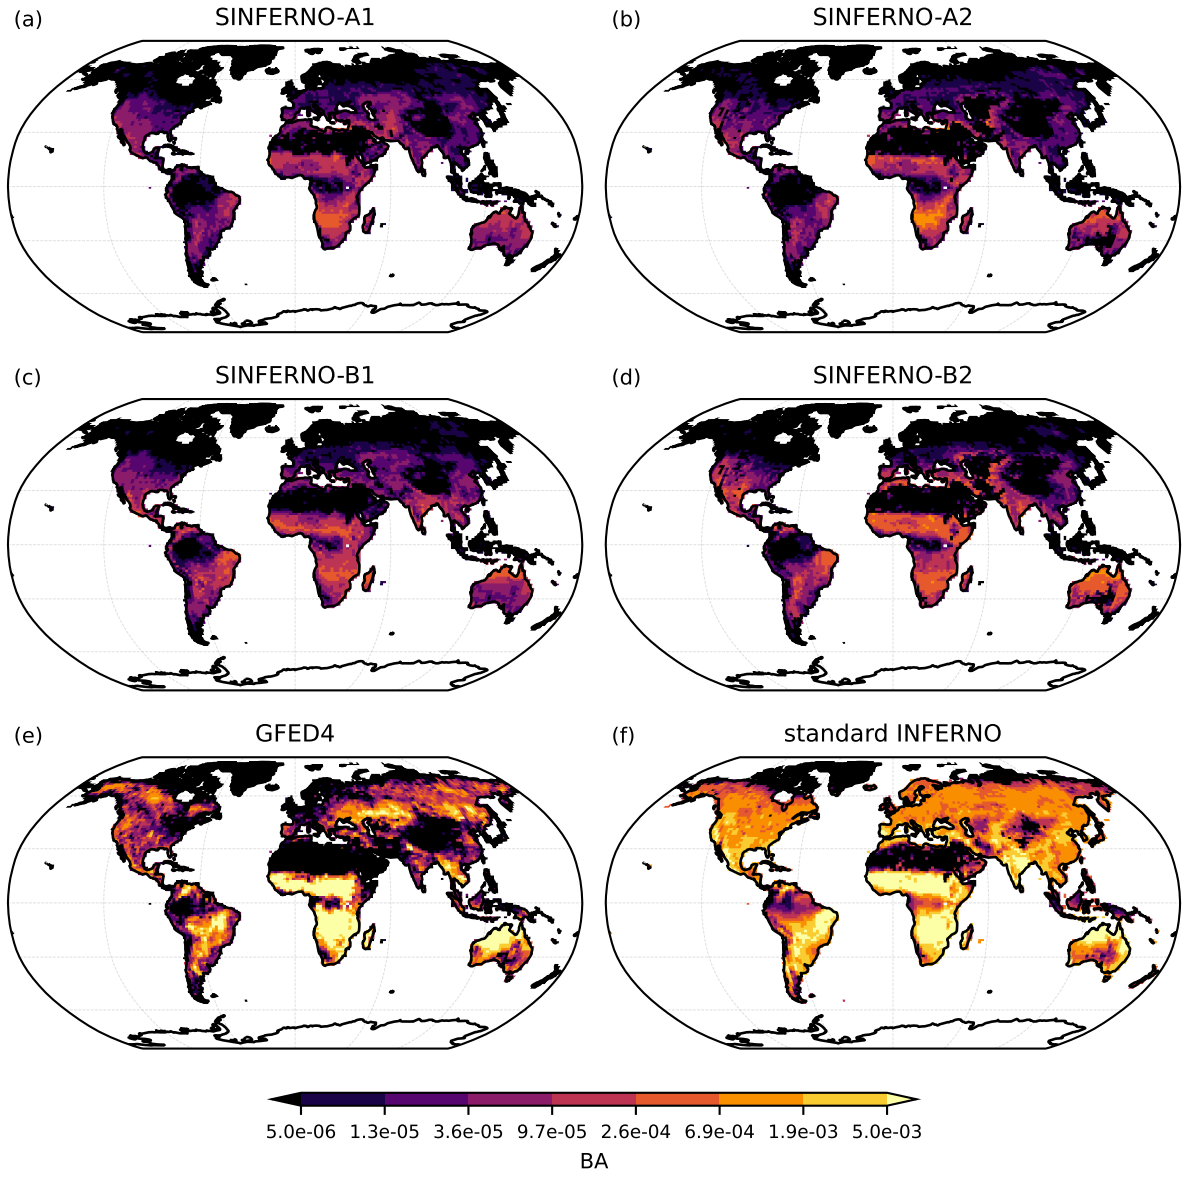


Figure 4.7: Modelled and observed BA on a log scale. The modelled BA by each parametrisation (a)–(d), the standard INFERNO results (f), and GFED4 observations (e) are shown.

4.3 Results

All parametrisations, including the standard INFERNO, adequately represent the major spatial patterns of burning (see Fig. 4.7 and 4.8). From Fig. 4.9 and 4.10, it is apparent that the B2 parametrisation has the best global performance as measured by the NME and the MPD scores. The MPD performance of the SINFERNO-B2 model was statistically significantly better than all other models but the SINFERNO-B1 model (see Fig. 4.10). However, while the SINFERNO-B2 model also had the best mean NME performance, the large spatial heterogeneity of the NME errors resulted in no statistically significant differences amongst the models (see Fig. 4.9).

Both SINFERNO models using the 'B' dryness approach based on VPD and precipitation performed best, however, outperforming the models using the 'A' dryness approach based on the dry day period. This difference is more pronounced for the phase errors compared to the magnitude errors, indicating that the combination of VPD and precipitation better accounts for the timing of fuel drying (or wetting).

The BA modelled by the SINFERNO models generally had a lower maximum magnitude compared to observations, with a maximum BA of around 1×10^{-3}. Observed GFED4 BA was generally below 1×10^{-2}, with only a few dozen months and grid cells exhibiting BA above this. Due to the skewed distribution of BA, most observations are below the 1×10^{-3} threshold. Furthermore, the standard INFERNO model is able to better represent these large BA magnitudes, as apparent from Fig. 4.7. Therefore, the failure of the SINFERNO models to accurately represent fire behaviour in high-BA regions likely contributes to their poor MPD performance compared to the standard INFERNO in regions such as grassland and dry shrubland (Fig. 4.11c) and savannah and dry woodland (Fig. 4.11d). Still, the SINFERNO-B2 parametrisation reproduces the spatial heterogeneity of BA the best overall, which is evident from its lowest NME score (see Fig. 4.9) and largely due to its better reproduction of low orders of magnitude of mean BA in boreal regions, for example.

While the phase performance of the B2 parametrisation globally outperforms the standard INFERNO parametrisation (see Fig. 4.10), this improvement is not universal, which is another

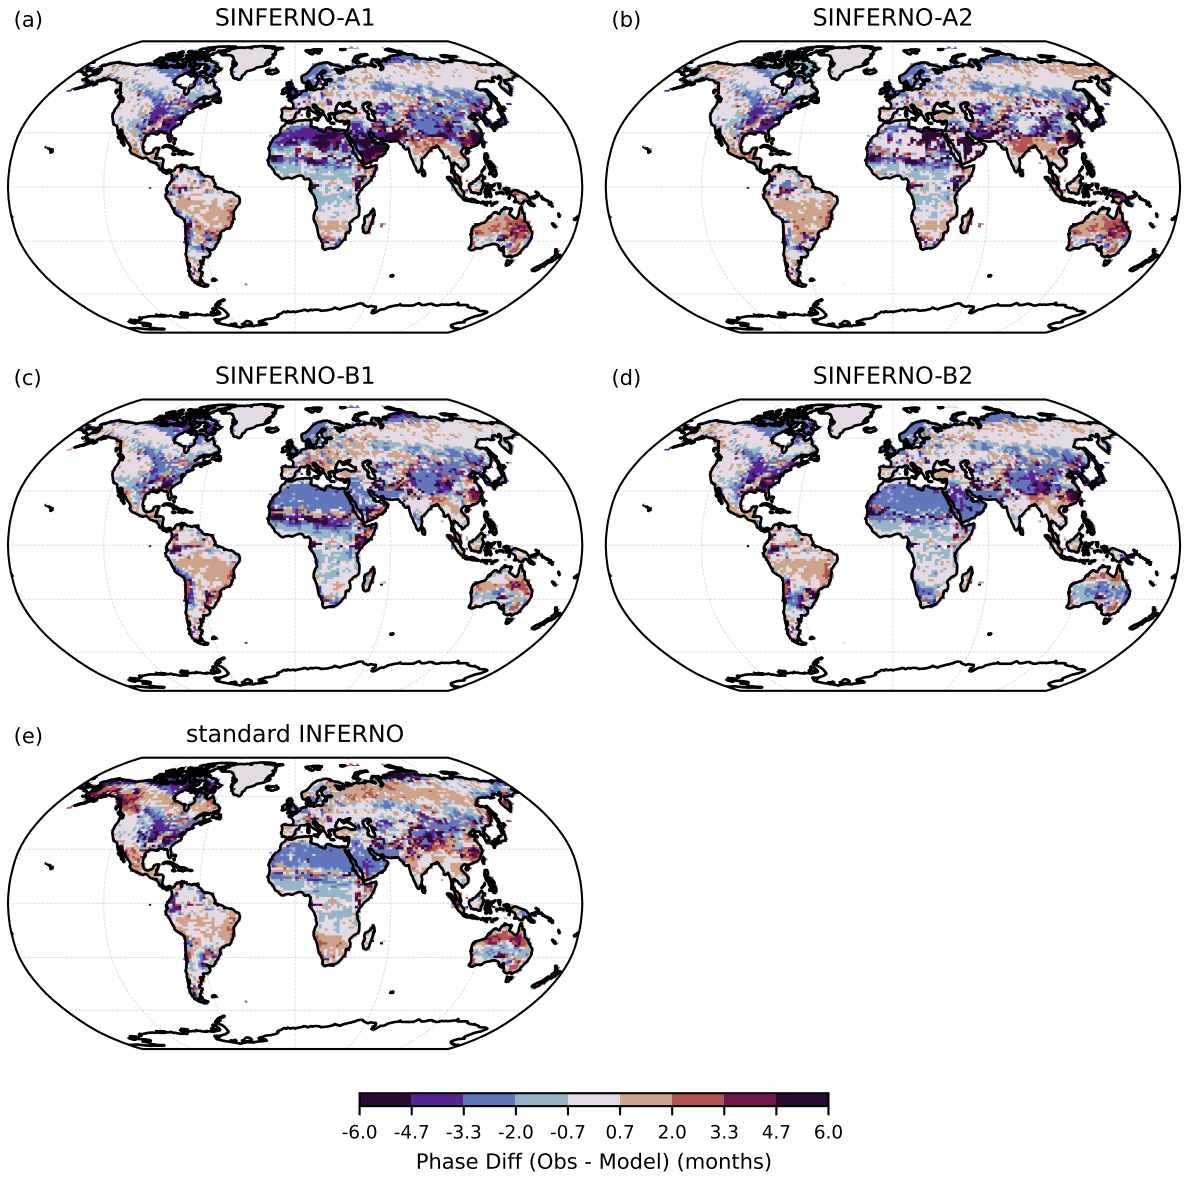


Figure 4.8: Phase differences between modelled and observed BA. The major spatial patterns of burning are adequately represented by all parametrisations, including the standard INFERNO. Comparing panels (d) and (e), it is apparent that the phase of burning has improved, especially in boreal regions (e.g. Alaska). SINFERNO-B2 has a phase error of around 1 month in this region, while the standard INFERNO predicted BA around 3 months too late. Similar improvements are also seen across the boreal regions and northern Australia, although phase performance did not improve in all regions, with China and the Cerrado being notable exceptions.

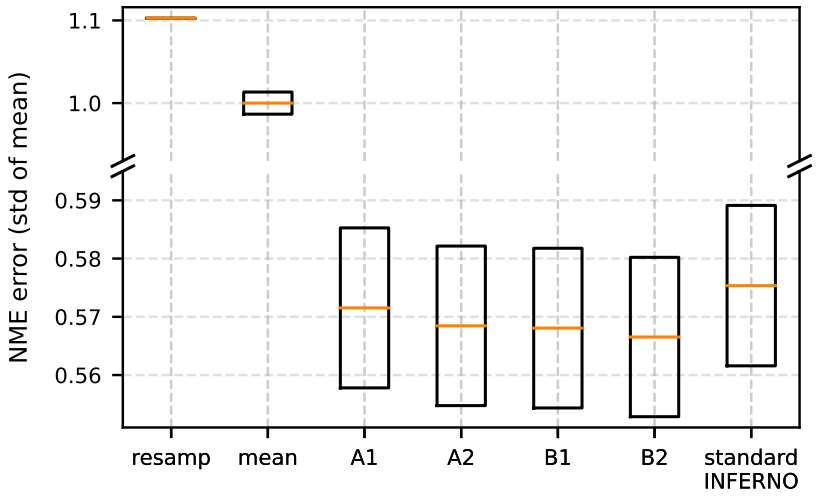


Figure 4.9: Average NME errors of the different parametrisations. Resampled and mean errors are derived from null-model analysis (see Section 4.2.8.1). Note the different scales above and below the y-axis break. With the best-performing B2 model exhibiting no significant difference compared to the standard INFERNO, $t(186502) = 0.45, p = 0.65$, it is clear that there are no significant differences between any of the SINFERNO models and the standard INFERNO.

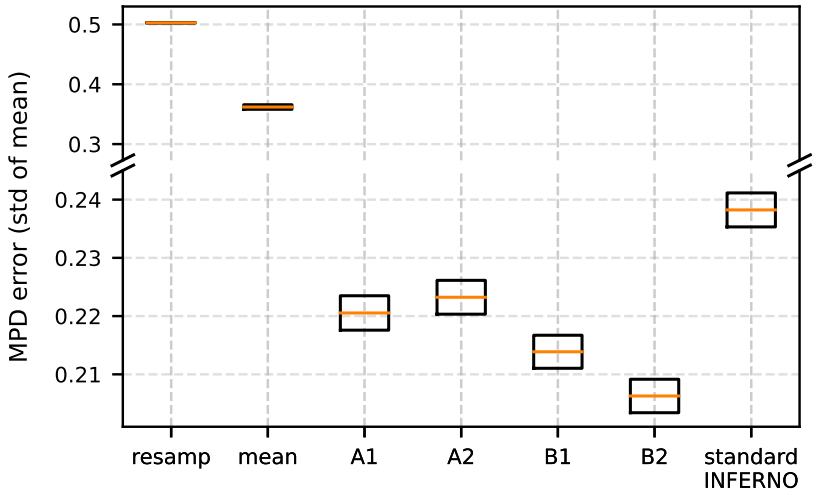


Figure 4.10: Average MPD errors of the different parametrisations. Resampled and mean errors are derived from null-model analysis (see Section 4.2.8.1). Note the different scales above and below the y-axis break. Using one-tailed t-tests, the following statistically significant MPD performance improvements were found: A1 < standard INFERNO, $t(11235) = -4.26, p < 0.001$; A2 < standard INFERNO, $t(11059) = -3.64, p < 0.001$; B1 < standard INFERNO, $t(11235) = -5.98, p < 0.001$; B2 < standard INFERNO, $t(11146) = -7.80, p < 0.001$; B2 < A1, $t(11149) = -3.46, p < 0.001$; B2 < A2, $t(10973) = -4.15, p < 0.001$, confirming that the B2 SINFERNO model performs significantly better than all but the B1 model, with all other SINFERNO models outperforming the standard INFERNO significantly.

reflection of the spatial heterogeneity of the performance improvements. As shown in Fig. 4.11, the phase performance of the different parametrisations is relatively close, except for the tundra and boreal forest regions, where the standard INFERNO performs noticeably worse. While the standard INFERNO outperforms the new SINFERNO models in the temperate and warm-temperate forest regions, the combined number of grid cells in these regions (1631) is much lower than the number of grid cells encompassed by the boreal forest and tundra regions (2818), where the SINFERNO models outperform the standard INFERNO.

Aside from Fig. 4.11, improvements to the phase performance made by the SINFERNO models are also made apparent by comparisons between Fig. 4.8d and 4.8e, where it can be seen that the phase of burning has improved, especially in boreal regions (e.g. Alaska). Whereas the standard INFERNO predicted BA around 3 months too late, the SINFERNO-B2 model reduces this error to around 1 month or less. Similar but less significant improvements are seen across the boreal and tundra biomes. Phase improvements were also made by the B2 model across northern Australia.

However, phase performance did not improve everywhere. For some regions affected by anthropogenic interventions, e.g. agricultural burning, such as the Cerrado or eastern China, phase errors generally maintained their sign (i.e. too early or late) while increasing in magnitude. Phase performance in certain other regions affected by anthropogenic interventions, such as the eastern US and Europe, remained comparable to the standard INFERNO.

4.4 Discussion

While all SINFERNO models and the standard INFERNO were able to capture the overall global pattern of burning, there were significant differences in their overall performance, with the phase performance, as measured by the MPD, indicating especially large improvements of the SINFERNO-B2 model over the standard INFERNO.

The metric that was optimised throughout was a composite loss comprised of arcsinh-NME and MPD. Note that the NME error takes into account both the magnitude and the timing of

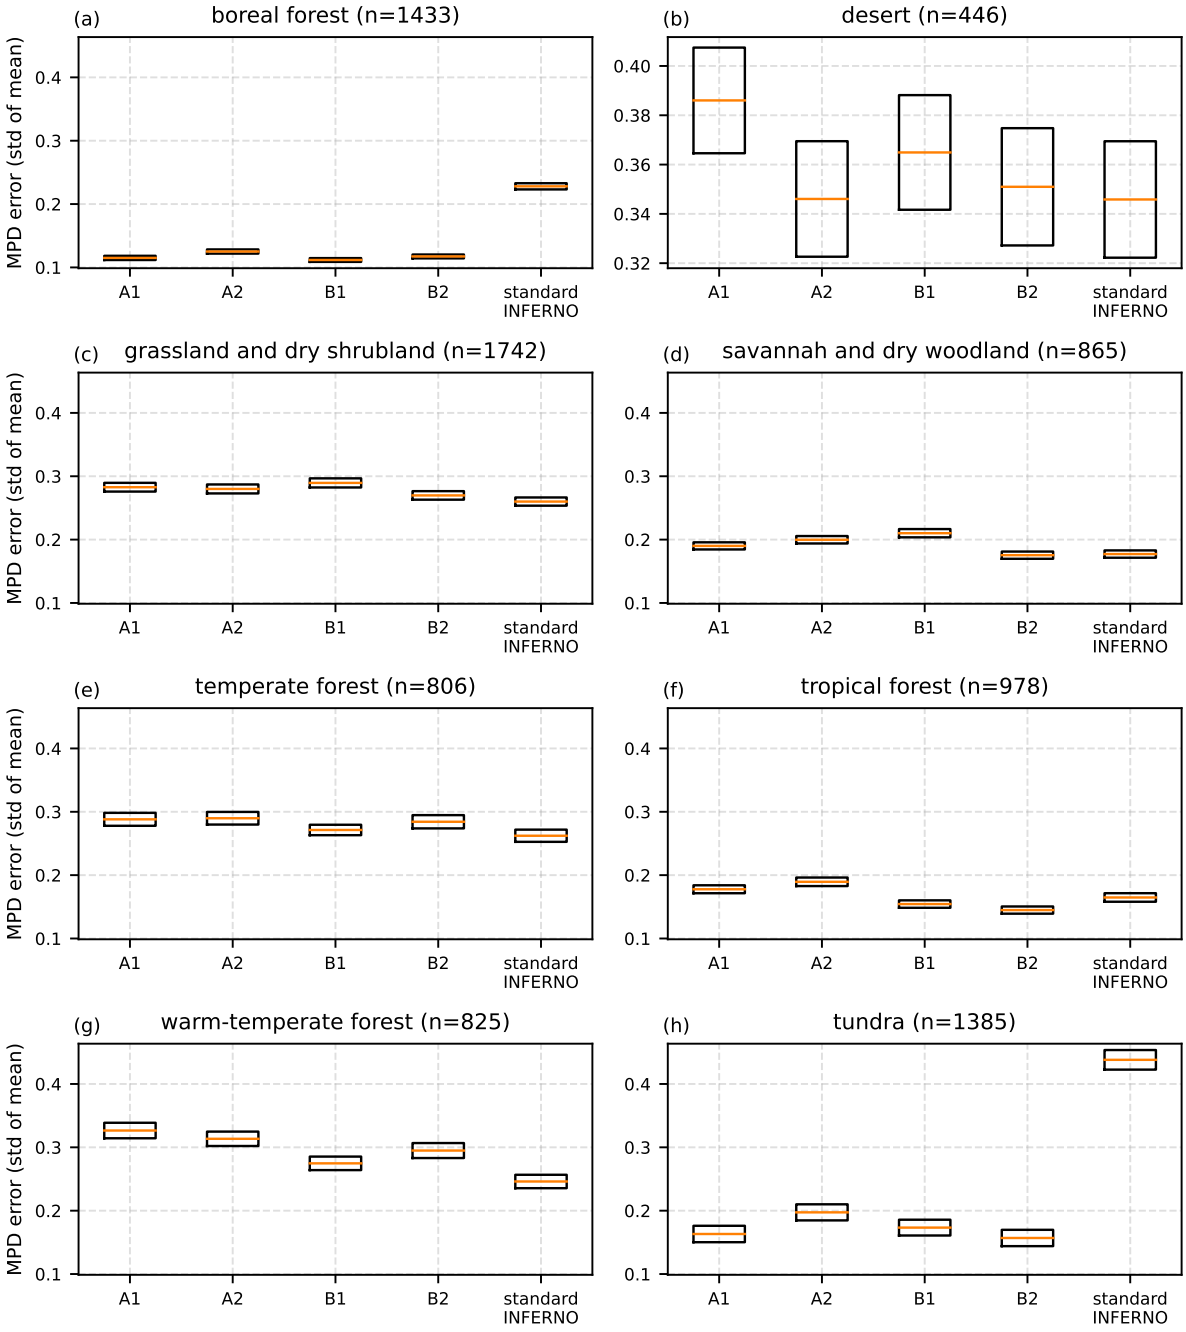


Figure 4.11: Regional MPD errors of the different parametrisations. The number of grid cells, n, is indicated for each PNV biome (see Fig. 4.6). All y-axes except for plot (b; desert) share the same limits. Most regions exhibit only small differences between the SINFERNO and the standard INFERNO models, but the boreal forest (a) and tundra (h) biomes show large performance gains by all the SINFERNO models compared to the standard INFERNO.

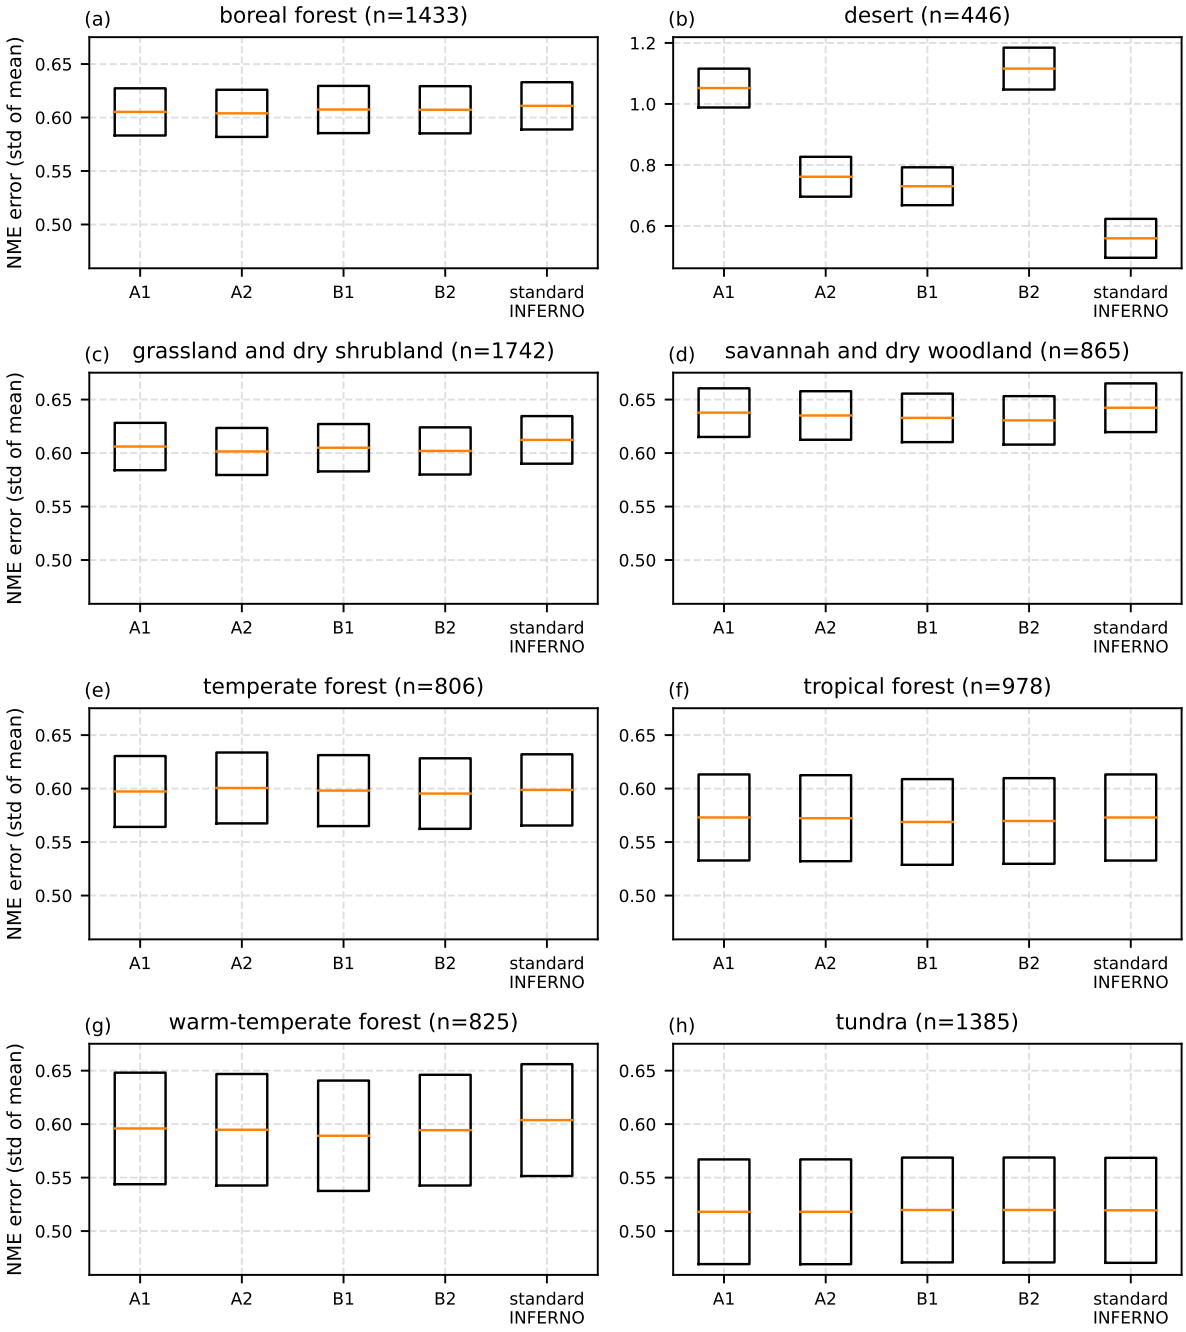


Figure 4.12: Regional NME errors of the different parametrisations. The number of grid cells, n, is indicated for each PNV biome (see Fig. 4.6). All y-axes except for plot (b; desert) share the same limits. It is apparent that for all key regions (ignoring (b)), there are no significant differences in NME performance between the SINFERNO and standard INFERNO models.

burning, while the MPD error, which is calculated on a per grid cell basis before being averaged across regions, only takes into account the timing, completely disregarding the mean annual BA at each grid cell. Compared to an optimisation of the untransformed NME, the arcsinh-NME placed emphasis on regions with low orders of magnitude of BA, which is reflected in the large MPD improvements seen when comparing the SINFERNO models to the standard INFERNO in boreal and tundra regions.

The reason why only small improvements in the NME are seen when comparing SINFERNO models to the standard INFERNO in Fig. 4.9 and 4.12, especially in the boreal forest and tundra biomes where improvements would be expected to be the largest, is because it is the untransformed NME error that is being plotted instead of the more unconventional transformed arcsinh-NME metric. Nonetheless, the MPD improvements in biomes with typically low BA demonstrate the effectiveness of the composite loss in training a model that is effective globally, as opposed to predominantly in regions with the highest BA. The standard INFERNO model was found to outperform the SINFERNO models in some biomes, predominantly biomes characterised by high BA, such as savannah and dry woodland. This is undoubtedly a consequence of optimising the arcsinh-NME and MPD composite loss as opposed to simply the NME or the sum of squares model-misfit since the composite loss places far less emphasis on capturing the few extremely large BA values that are observed.

It was found that using a combination of VPD and precipitation (dryness approach 'B') resulted in a better phase performance (as measured by the MPD) compared to the other SINFERNO models, likely due to the better representation of the timing of fuel drying (or wetting). This supports the findings of Drüke et al. [2019], who found that a VPD-based flammability parametrisation performed well. While SINFERNO-B2 was found to perform the best overall, it also has the most parameters, as both the dryness scheme and the fuel build-up scheme involve multiple variables that are subject to optimisation. For this reason, it may be expected to outperform the other parametrisations. The extent to which overfitting is responsible for the improved performance of the B2 parametrisation is therefore explored using

cross-validation (CV) in Chapter 5. However, the highly discontinuous, step-wise relationships fit by the B2 model (see Section 5.3.1) suggest that this increased performance may come at the cost of lower generalisability, indicating that which of the new parametrisations is best may depend on the desired application, particularly the degree of generalisability.

As noted by Knorr et al. [2014], there is increased temporal and spatial heterogeneity in the observations in areas with long fire return intervals (FRIs) (e.g. humid tropics or boreal forest biomes), with many grid cells typically having a particularly low BA. This is due to the stochasticity of mean fire activity in these regions, given the relatively short observational time coverage of my analysis compared to the FRI. Empirical models like the standard INFERNO and SINFERNO are, without the addition of any stochastic process, unable to predict anything but the expected mean fire activity (or fire risk) for a given location. This particular lack of accurate fire statistics is somewhat ameliorated by my use of climatological data across 17 years (2000–2016). The stochastic nature of fire occurrence in regions with a relatively large FRI nevertheless limits the achievable model performance while also suggesting that the BA modelled by e.g. SINFERNO should, by design, not approach the highest orders of magnitude in the observations over the limited available timespan. This is because it is the expected long-term mean fire activity that is being modelled instead of the actual fire occurrence. In the absence of BA observations over sufficiently long time periods relative to the FRI, this is a necessary trade-off.

The decrease in the predictive performance of the SINFERNO models compared to the standard INFERNO in some regions affected by significant anthropogenic intervention suggests that the use of cropland fraction (but ignoring population density) is not able to capture the anthropogenic impacts on fire. Note that since the cropland fraction typically does not have a large intra-annual variability, any related ignition (and suppression) behaviour was essentially constant on the climatological timescales investigated here. This supports the findings of Forkel et al. [2017], who found that including cropland as a measure of human influence did not significantly improve results. However, the seasonal distribution of fires is significantly affected by human

intervention, especially in regions used for crops and livestock [Le Page et al., 2010, Benali et al., 2017]. Therefore, usage of additional data on human impacts, such as the HDI used in Teixeira et al. [2021] or the even more detailed database of anthropogenic fire impacts (DAFI) [Perkins et al., 2021], in addition to cropland and pasture coverage, may be needed to adequately capture anthropogenic impacts on fire.

Since cropland cover, which represents the potential for anthropogenic ignitions and suppression, is often practically constant, and time-varying natural ignitions (e.g. lightning) are not considered here, the improvements made by the SINFERNO models to the phase of the predicted burning compared to the standard INFERNO in many regions are likely due to improved representation of climatic and vegetation influences, e.g. on fuel wetting and drying. And while the random forest (RF) models developed in Chapter 3 could not capture BA in boreal forests well, they were trained using an error metric that placed significant emphasis on regions with the highest BA (e.g. the Sahel). Nonetheless, the SINFERNO models are mostly able to match the performance of the standard INFERNO in biomes with high BA, like grassland and dry shrubland or savannah and dry woodland, while outperforming it in biomes with low BA, like boreal forest and tundra. This indicates that the explicit inclusion of antecedent conditions in the SINFERNO models permits improved BA modelling across biomes. It also suggests that the RF models from Chapter 3 could have performed better in boreal regions had a different error metric (or, similarly, a transformed BA target variable) been used.

4.5 Conclusions

The development of a new flammability parametrisation for the empirical INFERNO fire model generally yielded improved performance, especially in regions with typically low BA. On a global scale, there were improvements to both the magnitude (NME) and the phase (MPD) of the predicted BA. However, regionally, a trade-off was observed with the performance of the new SINFERNO models worsening in some regions with typically high BA compared to the standard INFERNO model. Performance, especially the phase performance, was improved significantly

in northern regions typically exhibiting low BA. Such regional discrepancies suggest that the overall performance of the new SINFERNO global empirical fire models may be limited by their inherently parsimonious approach, the limitations of the coarse spatial scales involved, and the accuracy of the productivity predicted by the underlying JULES dynamic global vegetation model (DGVM). While offline runs of the new SINFERNO fire models enabled the consideration of all grid cells with reasonable computing resources, their use also limited the achievable realism of the simulations compared to online runs. Since offline runs do not capture the impact of the simulated BA on vegetation mortality, their use in the optimisation and evaluation of the new parametrisation thereby also limited the attainable performance.

Chapter 5

Sensitivity Analysis

Abstract

Despite the ability of contemporary fire models to reproduce the general patterns of global burnt area (BA), their phase performance has been isolated as a particular area for improvement, with poor reproduction of the relationship between vegetation and flammability as a likely cause. The underlying relationships between the input variables and modelled flammability, as well as the significance of the different model input variables and parameters, need to be understood in order to improve an investigated model. Here, I investigated the underlying relationships of the Sigmoidal INFERNO (SINFERNO) model alongside the significance of its input data and parameters. Generalised additive models (GAMs) were used to verify the underlying sigmoidal relationships derived from the SINFERNO parameters. It was found that these were generally reasonable representations of the underlying relationships identified using the GAM empirical analysis. Sobol sensitivity analysis (SA) was used to identify the significance of the input data and model parameters, with dryness emerging as the most important variable, followed closely by temperature. The SINFERNO sigmoid centre parameter had the most significant influence on both the normalised mean error (NME) and mean phase difference (MPD) error metrics, although the remaining sigmoid parameters also predominantly influenced the NME, albeit to a lesser degree. Map plots of the Sobol SA significances revealed significant regional heterogeneity, with the temperature weight parameter being most important in the northern regions where the temperature has a dominant limiting effect on flammability, while the importance of the temperature centre parameter was more uniform, for example. Finally, SINFERNO models were constrained in their number of input data variables and parameters. Using stratified five-fold cross-validation (CV), this revealed that the ideal number of SINFERNO model parameters is around 30 compared to the possible 50 parameters.

5.1 Introduction

Having introduced a new class of flammability parametrisations based on the INteractive Fire and Emission algoRithm for Natural envirOnments (INFERNO), the Sigmoidal INFERNO (SINFERNO) model, in the previous chapter, the importance of its individual model components will be explored in this chapter. Using combinations of different dryness parametrisations (dry-days or vapour pressure deficit (VPD) and precipitation) and different fuel build-up parametrisations (antecedent net primary productivity (NPP) or a re-parametrised soil carbon pool), SINFERNO predicts burnt area (BA) using sigmoidal underlying relationships between these driving variables and flammability. The different combinations of driving variables from SINFERNO-A1 to SINFERNO-B2 (see Table 4.1) will be investigated here, although the main focus will lie on the best-performing SINFERNO-B2 model.

While SINFERNO employs a relatively parsimonious approach, this alone does not elucidate the significance of its individual components, which depends both on the parameters that result from the model optimisation and the data being fed into the model. For example, if flammability were to respond identically to temperature and dryness, the amount of variability in each of these two variables, as opposed to the underlying relationships between these variables and flammability, would directly determine the relative importance of the corresponding model components. In reality, both the model behaviour and the input are highly variable, prompting the use of further investigative tools to disentangle the relative importance of different aspects of the SINFERNO models.

SINFERNO models are optimised using a composite loss that emphasises areas with relatively low BA (via an arcsinh transformation of the BA) and the phase of the burning (via the mean phase difference (MPD) of the climatological predicted and observed BA). Chapter 4 showed that the nature of this composite loss impedes the evaluation of an analytical Jacobian, which could otherwise be used to aid model optimisation. However, this similarly impedes the analysis of the model components, which is the goal of this chapter. While the response of the phase error to different inputs is of interest, for example, the exact analytical expression linking all input

data and parameters to the global phase performance would be computationally prohibitive to evaluate. Therefore, this chapter will make use of techniques which can be used to estimate the overall global significance of the various SINFERNO components without requiring an analytical expression linking e.g. the MPD to the input driving variables.

5.1.1 Verification of the underlying relationships using GAMs

Compared to the random forest (RF) machine learning analyses presented in Chapter 3, the SINFERNO models employ a much more parsimonious approach. This means that the underlying sigmoidal relationships between the input data and flammability can be inferred from the model parameters alone without having to resort to more complicated analyses like accumulated local effects (ALE) plots, which rely on a large amount of data to make statistically robust claims about a model's average underlying relationships. Consequently, SINFERNO model behaviour (BA response to input data and parameter values) can be determined reliably at the grid-scale level for individual timesteps.

To evaluate the underlying sigmoidal relationships of the optimised SINFERNO models, these known relationships can be compared to those of a simple statistical model, e.g. a generalised additive model (GAM), fit to the same data as the SINFERNO models. For example, while it may appear intuitive that relationships between predictor variables and BA are thresholded, i.e. no longer increasing or decreasing after a certain point, establishing whether this emerges from the available data can be achieved by fitting GAMs (or other similarly flexible models) and inspecting the resulting underlying relationships. An advantage of GAMs in this context is that they, similarly to SINFERNO models themselves, provide explicit formulae for these underlying relationships as opposed to RFs, for example, where statistical methods such as ALEs would need to be applied to approximate them.

5.1.2 Global sensitivity analysis using Sobol SA

Despite the ability to determine the underlying sigmoidal relationships used by SINFERNO models from their optimised parameters, the sensitivity of the normalised mean error (NME)

and MPD error metrics to variations in the input data or parameter values remains practically inaccessible without further analyses. This is due to the complicated analytical expressions that link the error metrics to the input data and parameters. Global sensitivity analysis (SA) metrics like the Sobol metrics introduced in Section 5.2.2 can be used in order to estimate the significance of different input variables and parameters for such error metrics, however. Variance-based SA, or Sobol SA [Sobol′, 1993, 2001, Saltelli, 2002, Jacques et al., 2006, Saltelli et al., 2010, Owen, 2014, Norton, 2015], has been used in many different contexts, ranging from atmospheric science and traffic simulation to fire modelling [e.g. Gomez-Dans et al., 2018, Wild et al., 2020, Ge and Menendez, 2021]. And while simpler perturbations have previously been used for fire models to carry out SA, such as the fractional perturbations (i.e. +X%, -Y%) used in Mangeon [2017], Sobol SA yields much more robust estimates. This methodology (compared to simple one-at-a-time perturbations) takes into account a wide range of perturbations of multiple variables simultaneously. Thus, the model's global behaviour (i.e. as in across the entire range of input parameters or data, as opposed to simply around a single point in the input space of values) can be determined, albeit at a greater computational cost [Saltelli et al., 2008]. Here, Sobol SA is used to infer the sensitivity of the NME and MPD error metrics for individual grid cells, with the sensitivity across multiple grid cells being derived using averages of the individual sensitivities.

5.1.3 Reduced complexity SINFERNO models

Another way of determining which model parameters and data have the greatest impact on the BA modelled by SINFERNO is by limiting the parameters and data accessible to the model and evaluating which limitations have the most impact. Models limited in this fashion can be said to have a lower 'complexity' (i.e. fewer degrees of freedom), as there are fewer parameters to be tuned. To determine the SINFERNO model complexity (i.e. number of parameters and included predictors) required to achieve the ideal performance (generalisability), a parameter selection approach can be used. For example, variable and parameter selection was used by

Forkel et al. [2017] to select the best Satellite Observations to predict FIre Activity (SOFIA) model. The empirical SOFIA models have conceptual similarities to the SINFERNO approach. For the same data and parametrisation, Forkel et al. [2017] varied the number of parameters by switching between different land cover grouping schemes. Building upon their approach, I explore a more iterative, fine-grained approach here, where a series of models with varying complexities are constructed. Each model increases in complexity compared to the last, and cross-validation (CV) and the Akaike information criterion (AIC) are used in addition to the composite loss (arcsinh-NME and MPD) in order to investigate the ideal degree of complexity.

5.2 Methods

GAM analysis, Sobol SA, and reduced complexity modelling have different goals. While the GAM analysis investigates the shape of the underlying relationships, the Sobol SA and reduced complexity modelling both quantify the relative importance of different input data and parameters. The Sobol method analyses the response of the modelled BA to variations in the input data or parameters, while the reduced complexity models give insight into which parts of the model (i.e. parameters and the associated data these parameters act on) need to be included for maximum performance across a spectrum of model complexity (number of parameters).

5.2.1 GAM analysis

Verification of the underlying sigmoidal shape of the relationships was undertaken using GAMs instead of a machine learning model like the previously used RF model to avoid the additional uncertainties (e.g. due to extrapolation) that can accompany such an approach. The restricted maximum likelihood (REML) method was used to automatically select smoothing parameters, and a quasi-binomial distribution with a logit link function was used due to the highly skewed

distribution of the target BA variable. The formula used followed the following pattern:

$$\begin{aligned} y \sim\ & s(temperature_Trees, by = frac_Trees) \\ & + s(temperature_Grass, by = frac_Grass) \\ & + \cdots \\ & + s(dry_days_Shrubs, by = combined_frac_Shrubs). \end{aligned} \tag{5.1}$$

This emulates the approach taken by SINFERNO, where individual sigmoidal curves are used for each plant functional type (PFT)-group. The GAM models were fit to the same climatological data as the SINFERNO models, and the underlying relationships were extracted. As shown in Eq. (5.1), individual factors are computed for variables for each PFT-group (average variable across the PFT-group), weighted themselves by the cumulative PFT-group landcover at each grid cell. The sum of the resulting curves, acting upon the four groups of input data (e.g. temperature or dry days), is then fed into the inverse of the logit function, the sigmoid function. Contrary to an exponential link function, this is not equivalent to a multiplication of the individual curves for each variable since for the sigmoid function s, $s(a+b) \neq s(a) \cdot s(b)$. From this, it follows that the relationship between a and the output y depends on the other variables (with their collective influence being denoted b) when viewed in the untransformed output BA space. In other words, the effect of a given perturbation of a depends on the value of b. To determine the possible influence of each factor, the range of all data must therefore be taken into account.

5.2.2 Sobol sensitivity analysis

Sobol SA was used to determine the relative importance of either different data variables or different parameters for the BA predicted by the SINFERNO models. This is achieved by varying the considered parameters or input variables across a wide range of values and analysing the response of the modelled BA to these variations. The sensitivities thereby calculated using Sobol SA rely not only on the behaviour of the model being investigated but also on the variability of the tested variables or parameters. This variability (possible range of values) was

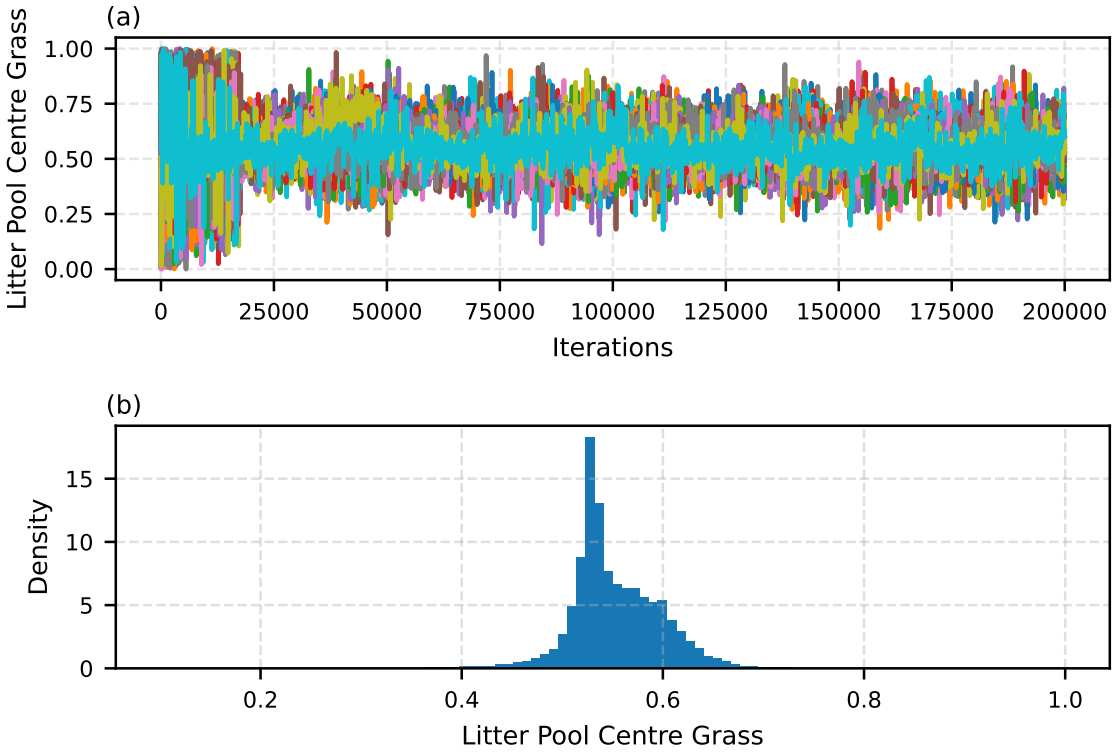


Figure 5.1: Convergence of DREAM MCMC chains for the `litter pool centre grass` parameter of the SINFERNO-B2 model. The last half of the chains are assumed to have reached convergence to the final (potentially multi-modal) distribution and therefore represent the likely range of the parameter values that match the data based on a given tolerance and the composite loss.

defined using lower and upper sampling bounds for this analysis. The sampling bounds were set differently for the data and parameters. Throughout this analysis, data and parameters were not considered simultaneously due to the fact that the Sobol method requires the input variables to be independent [Jacques et al., 2006, Saltelli et al., 2008, Chastaing et al., 2012, Wiederkehr, 2018, Ge and Menendez, 2021]. The standard deviation of the 17 mean yearly values (i.e. the inter-annual variability (IAV) of the yearly mean) at each gridpoint was used to set the data sampling bounds because this captures the expected scale of variability.

5.2.2.1 MCMC simulation

A Markov chain Monte Carlo (MCMC) simulation using the Differential Evolution Adaptive Metropolis (DREAM) algorithm [Vrugt et al., 2008, 2009, Sadegh and Vrugt, 2013, 2014, Vrugt, 2016] was used in order to establish parameter sampling bounds supported by the data in

an approximate Bayesian computation (ABC) framework [Marin et al., 2012]. The DREAM algorithm was run with the following parameters: `N` $= 2 \times 10^5$, `beta` $= 1$, `nChains` $= 50$, and `c` $= 0.1$. The acceptance 'delta' between the lowest previously found minimum using the basin-hopping approach ('1000,15' in Fig. 5.13) and the model parameters found by the DREAM MCMC sampler was set to 0.002. In practice, this meant that model parameters that performed worse than the previously found minimum by a margin of 0.002 (in composite loss units) were accepted and formed part of the chain history, as shown in Fig. 5.1a. An example of the convergence of multiple chains is shown in Fig. 5.1, with Fig. 5.1b showing the approximate distribution of the parameter values that met the chosen composite loss threshold. Note that the shape of the estimated parameter distributions varied, with some parameters being almost equally likely to be found across their entire allowed range for the MCMC analysis. For each parameter, the allowed parameter sampling bounds for the Sobol SA were determined using the mean of its sampled MCMC parameter distribution ± its standard deviation. The resulting sampling bounds were further constrained to be within the extrema observed during MCMC sampling.

5.2.2.2 Sobol SA grouping

The Sobol SA was then carried out using individual parameters grouped by data or parameter categories, both to avoid the analysis of dependent variables (for example, parameters for different PFT-groups are likely to be correlated) and to reduce the computational burden of the analyses. For example, the temperature sigmoid centre parameters across all PFTs were grouped together. A list of all parameter groups investigated for the SINFERNO-B2 model is shown in Table 5.3. Similarly, the investigated data groups are shown in Table 5.1.

Sobol SA with such groupings applies the same relative perturbation to each of the parameter or data variables within a given group. For example, consider the 'Temperature Centre' group, which consists of the parameters `temperature centre trees`, `temperature centre grass`, and `temperature centre shrubs` (one parameter per PFT-group). Note that the Sobol SA is

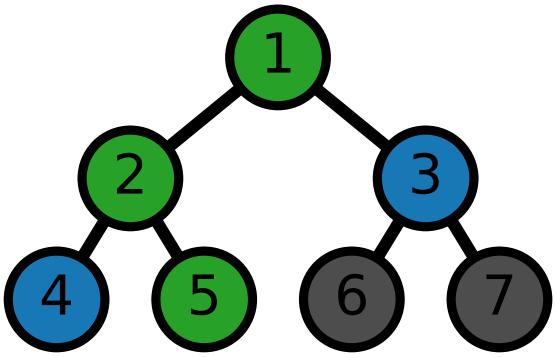


Figure 5.2: Iterative optimisation graph representation. Green nodes (1, 2, 5) are the chosen nodes representing the ideal performance amongst the explored nodes at each level. Explored but discarded nodes are shown in blue (3, 4). At each level, the best-performing model (green) is chosen as the basis for the models in the next level (i.e. $1 \rightarrow 2 \rightarrow 5$). Nodes branching off of green nodes are explored, but the grey nodes branching off of discarded nodes, like nodes 6 and 7 above, are never explored by the algorithm.

applied to multiple such groups in order to quantify the relative impact of each of these groups. If during the course of the Sobol SA, the `temperature centre grass` parameter is increased by half of its range, the `temperature centre trees` and `temperature centre shrubs` parameters will also be increased by half of their respective ranges. Thus, their significance can be calculated as a whole using far fewer model runs and avoiding the complications introduced by using the Sobol technique on dependent variables (see e.g. [Jacques et al., 2006, Saltelli et al., 2008, Chastaing et al., 2012, Wiederkehr, 2018, Ge and Menendez, 2021]). The initial gridpoint-specific significance values were also aggregated to overall global (geographically) significances using the mean of the individual significance values across grid cells. The standard deviation of the gridpoint-specific values was also calculated to provide an estimate of the significance variability.

In addition to calculating the first-order effect of a single variable, denoted S_1, the Sobol method also calculates the total effect S_T, which takes into account the first-order effects and any interactions with the other investigated variables. The sum of total effects across all variables, $\sum_i S_{T_i} \geq 1$, will only be equal to 1 for purely additive models, which do not have any interactions [Homma and Saltelli, 1996]. Comparing S_1 to S_T therefore provides an estimate of the degree to which interactions affect the model output.



Figure 5.3: Example of a relatively small, simple model configuration with many parameters being shared between PFT-groups as opposed to being individually optimised. A total of eight parameters are optimised in the configuration shown. All weights are either 0 or 1, meaning that they are not optimised. All parameters but the dryness-factor are shared across PFTs for either temperature and dryness separately. Other model categories are not present, so they do not inhibit burning at all. For temperature, there are three parameters (factor, centre, shape). For dryness, there are four parameters (two factors for different PFT-groups as indicated by the different-coloured blocks, centre, shape). The last remaining parameter (not shown) is the overall scaling factor which is always optimised.

5.2.3 Reduced complexity models

A 'reduced complexity' SINFERNO model is similar to a 'full' SINFERNO model, albeit with a reduced number of degrees of freedom, which is achieved by applying special parameter values and constraints to the full SINFERNO model. Reduced complexity SINFERNO models describe a subset of SINFERNO models where certain values are either set to 0 (e.g. a weight set to 0 implies ignoring the corresponding data variable(s)) or the effective number of degrees of freedom is otherwise reduced by mandating that certain parameters are set to the same value across certain PFT-groups. The latter refers to varying the way parameters are grouped across PFT-groups (the 'grouping scheme'). The grouping scheme employed here is more flexible than the one used by Forkel et al. [2017] for their SOFIA models since individual SINFERNO sigmoid parameters (e.g. the temperature-centre) may be shared across all three PFT-groups or distributed amongst subsets thereof (e.g. one value being jointly optimised for two of the PFT-groups, with another value being optimised for the remaining PFT-group). Overall, optimising the number of parameters of such reduced complexity SINFERNO models considers both the selection of data variables and parameters simultaneously.

This optimisation can be visualised as a graph search, illustrated here in Fig. 5.2, with different nodes in the graph representing different SINFERNO model configurations. A model

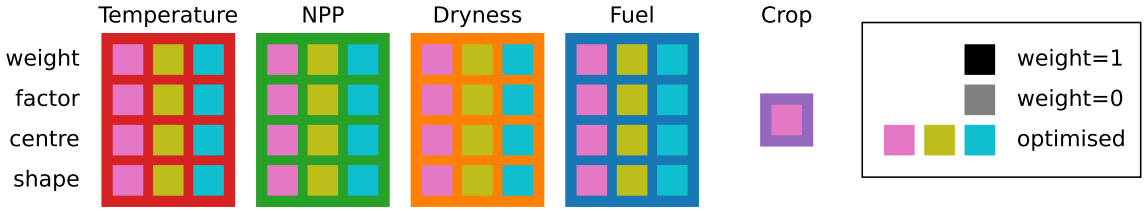


Figure 5.4: The most complicated model configuration possible, with no parameters shared across PFT-groups. A total of 50 parameters are optimised in the configuration shown. All model categories are present, with temperature, NPP, dryness, and fuel requiring 12 parameters each (48 in total). The weight, factor, centre, and shape parameters are optimised separately for each of the three PFT-groups. This is the case for each of the four categories. Cropland requires an additional parameter. The last remaining parameter (not shown) is the overall scaling factor which is always optimised.

configuration is the collection of possible model parameters that control the four different, potentially PFT-group-dependent aspects of the model and potentially also the influence of crop cover. A subset of possible configurations is explored by finding the best-performing configuration for a given complexity (number of parameters) before continuing the search at an increased complexity starting from this best-performing configuration. There is no backtracking, meaning that only configurations with higher complexity than previously explored configurations are considered as new configurations. While backtracking would be necessary to escape potential local minima, avoidance of backtracking was a compromise necessitated by the computational cost associated with optimising parameters for individual model configurations, which has to be repeated numerous times during this procedure. However, the 'overall' optimisation (exploration of different model complexities) was carried out twice using different settings for the basin-hopping and local minimisation algorithms used to optimise the reduced complexity SINFERNO models at each step of the overall optimisation. These settings were chosen to be less computationally demanding than those used in Chapter 4, where 1000 local minimiser iterations and a minimum of 15 basin-hopping iterations yielding the same minimum were used to optimise the SINFERNO models. The two settings used here were either 1000 local minimiser iterations and a minimum of 5 basin-hopping iterations or 50 local iterations and 1 basin-hopping iteration. Two different settings were used to evaluate the impact of the minimisation settings on the performance of the

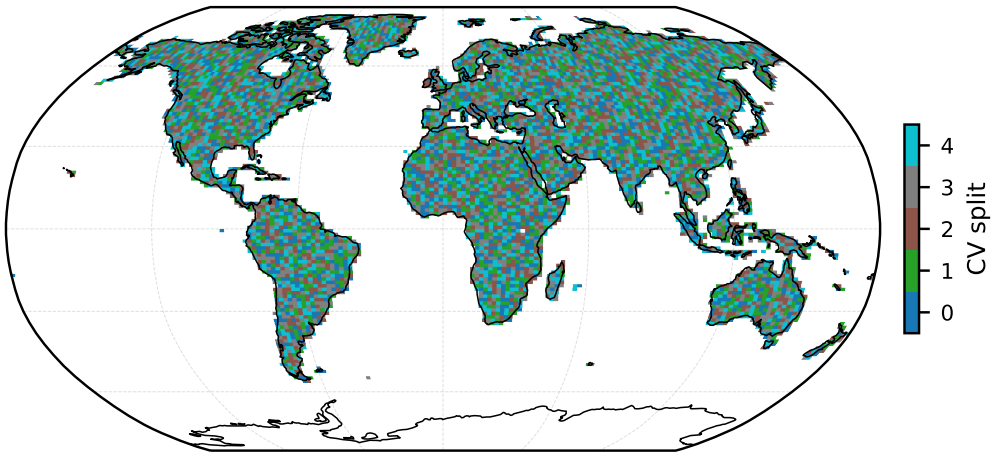


Figure 5.5: Geographical distribution of stratified five-fold CV splits with non-overlapping groups, distributed over the 7771 land points that comprise the JULES N96e grid.

reduced complexity SINFERNO models.

Conceptually, moving between configurations $2 \rightarrow 5$, as shown in Fig. 5.2, could be interpreted as moving between the configurations shown in Fig. 5.3 and 5.4, albeit with many intermediate steps omitted in Fig. 5.2. The actual graph of model states has a branching factor of around 40 initially, which shrinks down significantly towards the end of the algorithm as the model complexity reaches its limit of 50 parameters. This programmatic approach is highly flexible, exploring a wide variety of possible data inputs and land cover grouping schemes. But it does not include optimisation of the discrete parameters, the values of which were adopted from the previous, full optimisation. Each step is very fine-grained and could consist of the addition of a new set of parameters to capture the influence of, say, temperature (if not previously included) or the optimisation of the NPP sigmoid centre parameter separately for either grasses, shrubs, or trees (thereby effectively changing the landcover grouping scheme on a parameter-by-parameter basis).

For every stage of this iterative model fitting, i.e. for the best SINFERNO model for each number of optimised parameters, multiple metrics were computed to evaluate the ideal number of parameters. In addition to the performance (composite loss) itself, the AIC was employed, as in Forkel et al. [2017], to consider model performance in terms of the sum of squares error (SSE) while penalising models with too many parameters (both discrete and continuous). It was

calculated both with and without arcsinh transformation of the observed and predicted BA. Additionally, stratified five-fold CV, based on the composite loss, with non-overlapping groups was employed with the following BA bins: $(0, 1 \times 10^{-4}, 1 \times 10^{-2}, 0.1, 1)$ and using the `scikit-learn` `StratifiedGroupKFold` operator. Given the relatively low number of grid cells used here compared to the number of grid cells in Chapter 3, this approach was used instead of simple random five-fold CV in order to prevent some splits from containing far more grid cells with e.g. very high BA than others. Using the resulting splits (see Fig. 5.5), CV was used to quantify the generalisability (i.e. degree of overfitting) of the models across all stages of complexity. Both the CV results and AIC values were used to quantify the ideal model complexity.

5.3 Results

5.3.1 GAM analysis

The relationships fit by the GAM for temperature exhibit a similar shape to the sigmoidal relationships fit by the SINFERNO-B2 model: an initial steep rise followed by a plateau, as shown in Fig. 5.6a, e, i, and m.

The GAM model generally fit the expected negative relationship between instantaneous NPP and BA (Fig. 5.6f, j, and n). However, only the SINFERNO-B2 NPP relationship for shrubs agrees with the GAM relationships. For the Tree and Grass PFT-groups, the SINFERNO-B2 model fit flat, i.e. constant, relationships, which only agree with the relatively flat relationships fit by the GAM for low NPP.

There is a similar disagreement between the relationships fit by SINFERNO-B2 and the GAM for the Dry Bal variable, which represents the combined influence of VPD and precipitation. While the two models agree on the decreasing underlying relationship between Dry Bal and flammability for the Grass PFT-group, the SINFERNO-B2 model fit constant curves for both the Tree and Shrub PFT-groups, while the GAM fit virtually identical monotonically decreasing relationships between Dry Bal and BA for all three PFT-groups.

For the litter pool variable, SINFERNO-B2 fit steeply rising, almost step-wise relationships

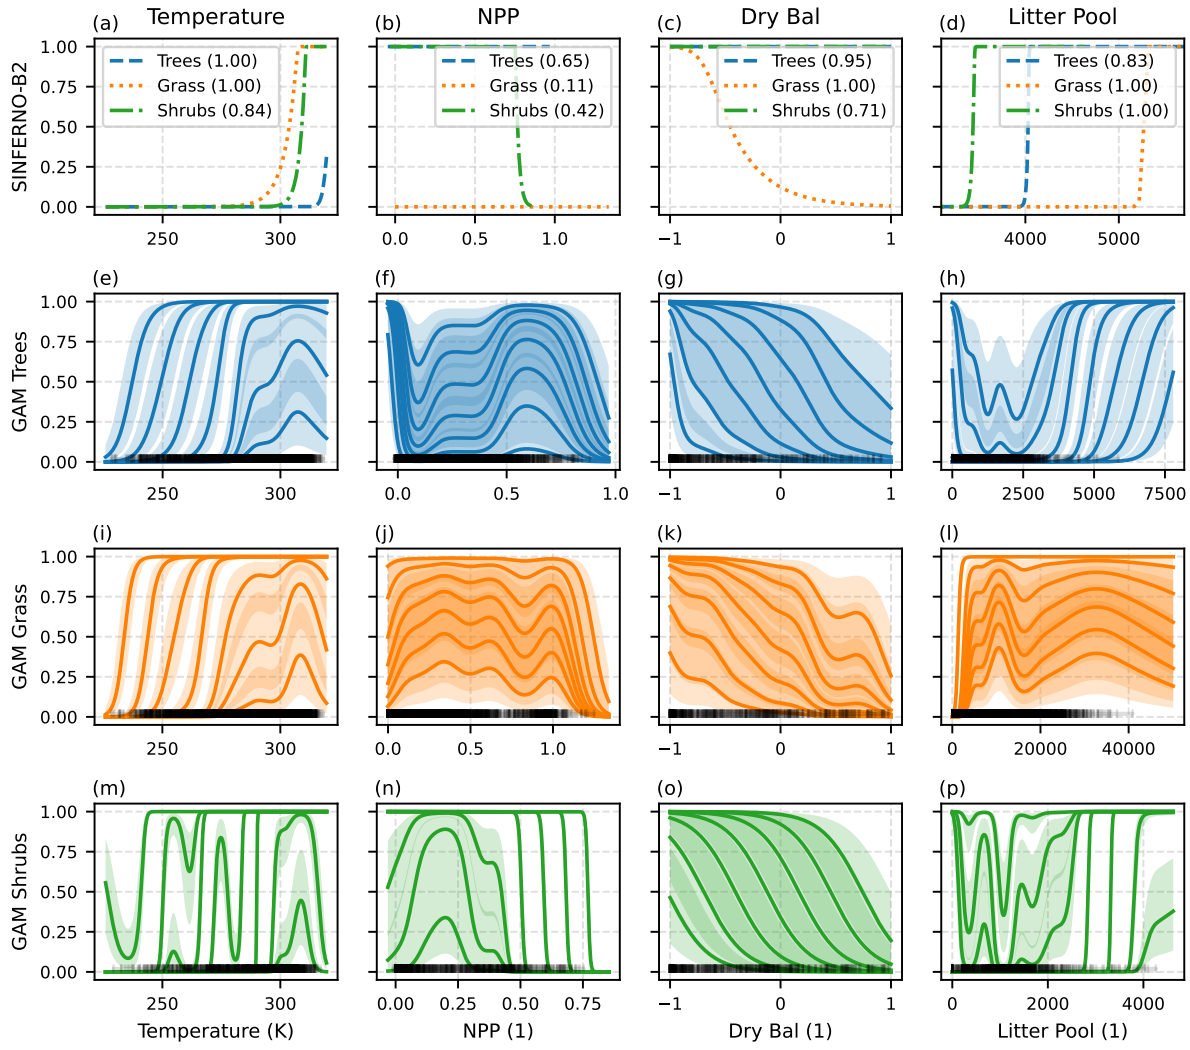


Figure 5.6: GAM verification of the underlying relationships in the SINFERNO-B2 model. Plots (a)–(d) show the relationships fit by SINFERNO-B2. Plots (e)–(p) show the relationships fit by the GAM. Due to the logit link used in the GAM, the individual underlying relationships vary depending on the values of the other data and are thus repeatedly plotted across a relevant range of the combined influence of these other data. Plotted around the GAM response curves are scaled standard errors. Rug plots for plots (e)–(p) indicate the distribution of the data for each PFT-group, with a darker shade denoting a greater abundance of data. The weights attached to each of the PFT-group sigmoids for the SINFERNO-B2 model are shown in brackets in the legends for plots (a)–(d). Note the different x-axis extents along each column, especially plot (d), which focuses on the range of interest with relationships in the omitted range being constant.

for all PFT-groups. These relationships essentially act as 'switches', either allowing fire or not. The impact of these switches is modulated by the weight parameter, which was 1 for both the Grass and Shrub PFT-groups, meaning that fire is completely suppressed when the litter pool is below the threshold. For the Tree PFT-group, the weight was 0.83, indicating that fire is only partially suppressed when the litter pool is below its threshold (which is around 4000).

Table 5.1: Global arcsinh-NME Sobol SA for SINFERNO-B2 input data variables. Both first-order effects (S_1) and total effects (S_T) are calculated. Mean effects (S_1, S_T) and corresponding standard deviations (S_1 std, S_T std) are calculated from individual effects at 7243 grid cells. For both S_1 and S_T, standard deviations can be seen to approach and often exceed the mean effects themselves, indicating that there is large regional heterogeneity.

	S_1	S_1 std	S_T	S_T std
Dry Bal	0.532 457	0.385 241	0.543 727	0.383 366
Temperature	0.282 087	0.315 358	0.303 176	0.323 469
Litter Pool	0.162 271	0.342 968	0.176 044	0.359 928
Cropland	0.000 304	0.004 037	0.000 301	0.004 031
NPP	0.000 000	0.000 014	0.000 000	0.000 008

On the other hand, the litter pool relationships fit by the GAM model are more heterogeneous when comparing the different PFT-groups. The GAM Grass PFT-group relationship corresponds most closely to its SINFERNO-B2 counterpart, increasing at around a value of 5000 (note the different x-axis extents in Fig. 5.6d and l). While the GAM Tree PFT-group relationship also increases at a point roughly matching its SINFERNO-B2 counterpart (∼4000, note also the different x-axis extents in Fig. 5.6d, h, l, and p), the GAM also fit a decreasing relationship with the Tree litter pool variable at low litter pool values, contrary to the monotonically increasing relationship fit by the SINFERNO-B2 model across all PFT-groups. For the Shrub PFT-group, the GAM relationship similarly reaches a threshold at a value of ∼3500, in agreement with SINFERNO-B2, while also fitting an oscillatory non-monotonic relationship below this threshold, thereby contradicting the SINFERNO-B2 model.

5.3.2 Sobol sensitivity analysis

The Dry Bal (dryness) and temperature variables are the most important input variables, as they have the highest ranking for both arcsinh-NME and MPD indices (see Tables 5.1 and 5.2, respectively). While the litter pool input variable (which represents fuel build-up) is the third most important variable for both arcsinh-NME and MPD, the standard deviation of its effect far exceeds its average total effect. Instantaneous productivity (NPP) and cropland cover are insignificant for both performance metrics. In general, the significance standard deviations often

Table 5.2: Global MPD Sobol SA for SINFERNO-B2 input data variables. Both first-order effects (S_1) and total effects (S_T) are calculated. Mean effects (S_1, S_T) and corresponding standard deviations (S_1 std, S_T std) are calculated from individual effects at 5342 grid cells. For both S_1 and S_T, standard deviations can be seen to approach and often exceed the mean effects themselves, indicating that there is large regional heterogeneity.

	S_1	S_1 std	S_T	S_T std
Dry Bal	0.463 677	0.443 667	0.469 019	0.443 541
Temperature	0.404 917	0.433 001	0.415 445	0.433 299
Litter Pool	0.120 362	0.298 740	0.126 855	0.308 516
NPP	0.000 036	0.002 222	0.000 041	0.002 688
Cropland	0.000 004	0.000 100	0.000 004	0.000 102

Table 5.3: Global arcsinh-NME Sobol SA for SINFERNO-B2 parameters. Both first-order effects (S_1) and total effects (S_T) are calculated. Mean effects (S_1, S_T) and corresponding standard deviations (S_1 std, S_T std) are calculated from individual effects at 7243 grid cells. For both S_1 and S_T, standard deviations can be seen to approach and often exceed the mean effects themselves, indicating that there is large regional heterogeneity.

	S_1	S_1 std	S_T	S_T std
Dry Bal Centre	0.147 390	0.126 842	0.254 075	0.168 535
Temperature Centre	0.137 531	0.143 586	0.218 592	0.181 713
Fuel Weight	0.100 380	0.116 462	0.169 228	0.165 485
Temperature Factor	0.068 326	0.109 482	0.131 229	0.172 023
Dry Bal Shape	0.089 706	0.165 687	0.112 984	0.175 627
NPP Centre	0.063 475	0.071 437	0.112 778	0.132 423
NPP Weight	0.054 304	0.064 096	0.103 055	0.086 585
Temperature Weight	0.044 909	0.068 190	0.078 013	0.108 156
Dryness Weight	0.026 767	0.051 006	0.053 844	0.070 842
Dry Bal Factor	0.007 992	0.017 689	0.024 364	0.034 277
Litter Pool Centre	0.012 768	0.075 670	0.024 204	0.110 141
Litter Pool Factor	0.006 727	0.044 583	0.014 597	0.088 160
Temperature Shape	0.001 915	0.006 340	0.008 045	0.020 074
NPP Factor	0.000 180	0.000 374	0.000 373	0.000 472
Litter Pool Shape	0.000 004	0.000 060	0.000 021	0.000 168
NPP Shape	0.000 004	0.000 031	0.000 002	0.000 003

Table 5.4: Global MPD Sobol SA for SINFERNO-B2 parameters. Both first-order effects (S_1) and total effects (S_T) are calculated. Mean effects (S_1, S_T) and corresponding standard deviations (S_1 std, S_T std) are calculated from individual effects at 5536 grid cells. For both S_1 and S_T, standard deviations can be seen to approach and often exceed the mean effects themselves, indicating that there is large regional heterogeneity. Note that small negative values can result from floating-point errors.

	S_1	S_1 std	S_T	S_T std
Dry Bal Centre	0.210 607	0.177 827	0.477 817	0.232 119
Temperature Centre	0.171 919	0.189 397	0.456 726	0.200 139
NPP Centre	0.077 536	0.087 947	0.229 548	0.201 774
Temperature Factor	0.026 937	0.038 498	0.184 740	0.154 526
Dry Bal Shape	0.018 432	0.027 274	0.088 299	0.092 653
NPP Weight	0.002 587	0.007 252	0.070 165	0.052 438
Temperature Weight	0.007 854	0.014 354	0.066 599	0.081 292
Fuel Weight	0.005 842	0.026 657	0.066 250	0.069 174
Dry Bal Factor	0.006 820	0.012 856	0.064 234	0.074 205
Temperature Shape	0.003 668	0.015 703	0.064 219	0.120 751
Dryness Weight	0.000 811	0.004 604	0.039 043	0.041 975
Litter Pool Centre	0.015 546	0.082 593	0.032 983	0.135 811
Litter Pool Factor	0.002 463	0.019 665	0.012 408	0.056 761
NPP Factor	0.000 035	0.001 737	0.006 794	0.009 868
NPP Shape	−0.000 009	0.000 623	0.000 602	0.002 265
Litter Pool Shape	0.000 001	0.000 532	0.000 539	0.005 470

approach, if not exceed, the mean global importances. Interactions, as indicated by the difference between S_T and S_1, are insignificant for all input predictor variables.

The dryness sigmoid centre parameter was consistently the most important parameter for both the arcsinh-NME and MPD indices (Tables 5.3 and 5.4, respectively). The temperature centre was also consistently important, supporting the finding that the temperature and dryness predictor variables are the most significant. For instantaneous NPP, the centre parameter was the most important sigmoid parameter. In general, the centre and weight parameters were most important (see Table 5.5). The centre parameters were assigned the highest proportion of significance using both the arcsinh-NME and MPD indices. However, the MPD-derived Sobol indices assigned much greater importance to the centres of the sigmoidal relationships: the importance proportion of the centre variables is 64 % using the MPD metric, compared to 47 %

Table 5.5: Proportion of mean global parameter importance by type measured using Sobol S_T indices for SINFERNO-B2. Arcsinh-NME SA indices were computed from individual significance values at 7243 out of 7771 land points. MPD SA indices were computed from individual significance values at 5536 out of 7771 land points. Proportions of the S_T indices were calculated using the mean global indices for each type of parameter (e.g. 'Centre' parameters) shown in Tables 5.3 and 5.4.

	arcsinh-NME		MPD	
	Proportion	Proportion std	Proportion	Proportion std
Centre	0.467 019	0.462 405	0.643 254	0.420 502
Weight	0.309 590	0.348 020	0.130 070	0.135 537
Factor	0.130 660	0.300 771	0.144 106	0.194 353
Shape	0.092 731	0.270 829	0.082 570	0.163 697

for the arcsinh-NME metric (see Table 5.5).

For controlling the NME, i.e. the magnitude of BA, the dryness variable is seen to be very important globally (Fig. 5.7b). The importance of temperature (Fig. 5.7a) is clustered more sporadically but is notably high in northern regions, e.g. Alaska, which saw major improvements compared to the standard INFERNO. Similarly, although the fuel build-up variable is not important on a global scale (Fig. 5.7c), it is again seen to play a significant role in the northern regions that perform much better in the SINFERNO models compared to the standard INFERNO.

For the MPD, the temperature is seen to have a significant effect in northern regions (Fig. 5.8a), while dryness predominantly affects the remainder of the globe (Fig. 5.8b). Again, fuel build-up (Fig. 5.8c) affects several disparate regions, including the northern regions that were seen to perform significantly better using the new parametrisation.

Overall, the litter pool fuel build-up parametrisation seems to respond well to the seasonality of antecedent productivity, thereby leading to improved phase performance in regions with high seasonality. Although both the dryness and temperature variables are significant on a global scale, the temperature has a greater influence on phase performance in northern regions. Both variables contribute to the NME magnitude performance in northern regions, however. For the northern regions, this suggests that the dryness parametrisation acts primarily to scale the magnitude of burning, while the temperature parametrisation affects both its magnitude and

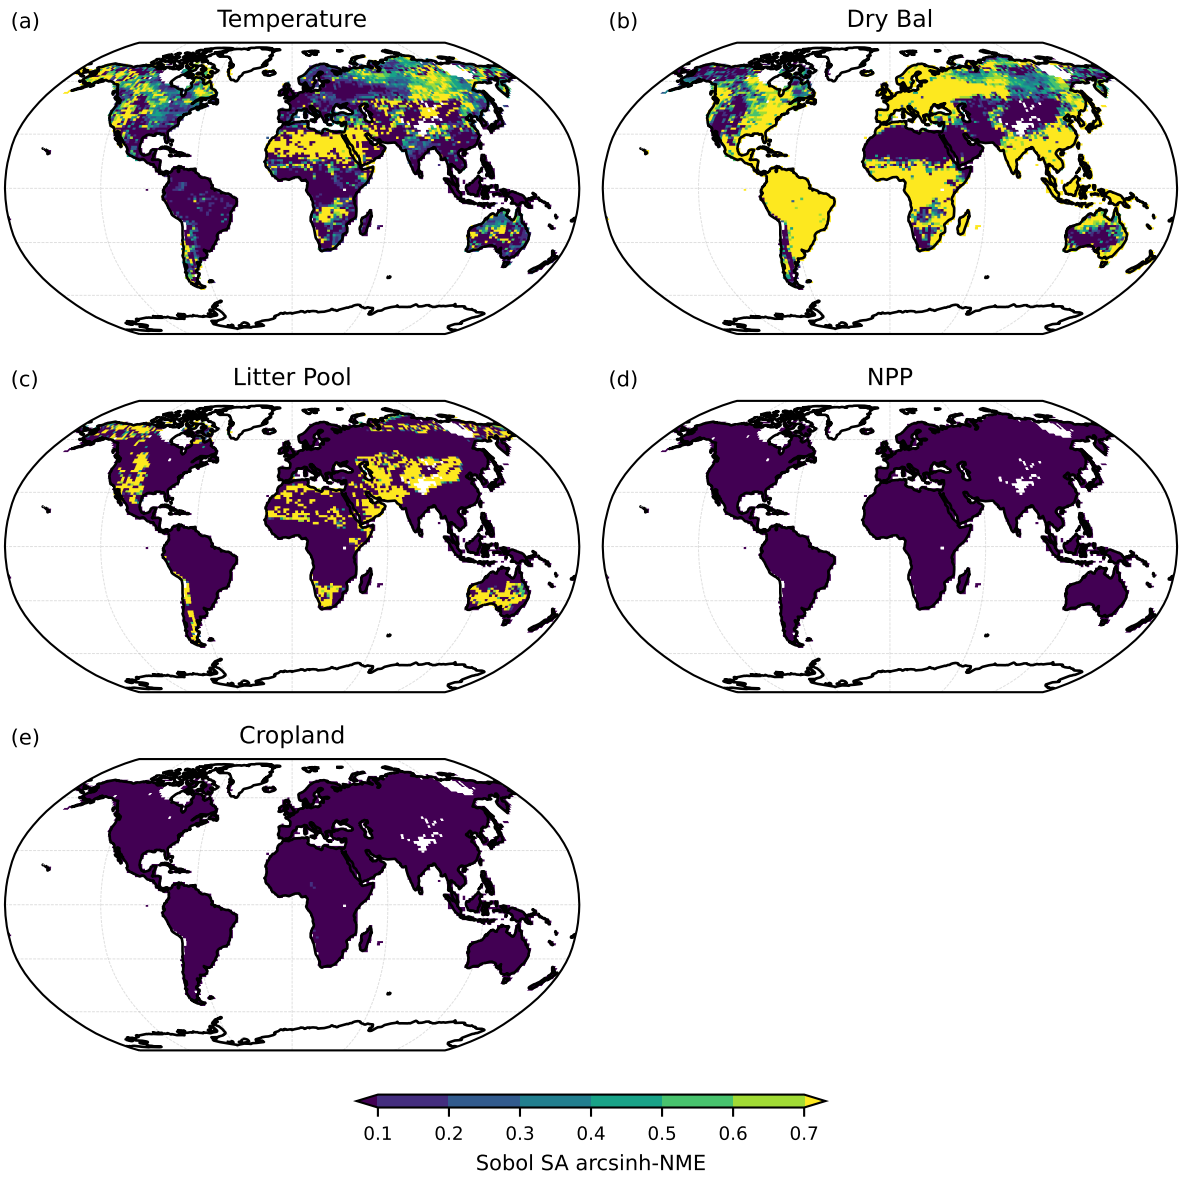


Figure 5.7: Sobol arcsinh-NME S_T importances for SINFERNO-B2 data. Importances were calculated at 7243 grid cells out of 7771 land grid cells. Dryness (Dry Bal) (b) is the most important globally. Temperature (a), while less significant overall, is generally important in regions where dryness has relatively low importance, such as some northern regions (e.g. Alaska). Fuel build-up (Litter Pool) (c) is generally important in regions where dryness is not important. Instantaneous productivity (NPP) (d) and cropland (e) are insignificant.

timing.

The instantaneous productivity and cropland data variables appear not to have a significant influence on NME (Fig. 5.7) and MPD (Fig. 5.8). However, Tables 5.3 and 5.4 show that parameters related to the instantaneous NPP have a noteworthy impact despite the low significance of the corresponding data variable.

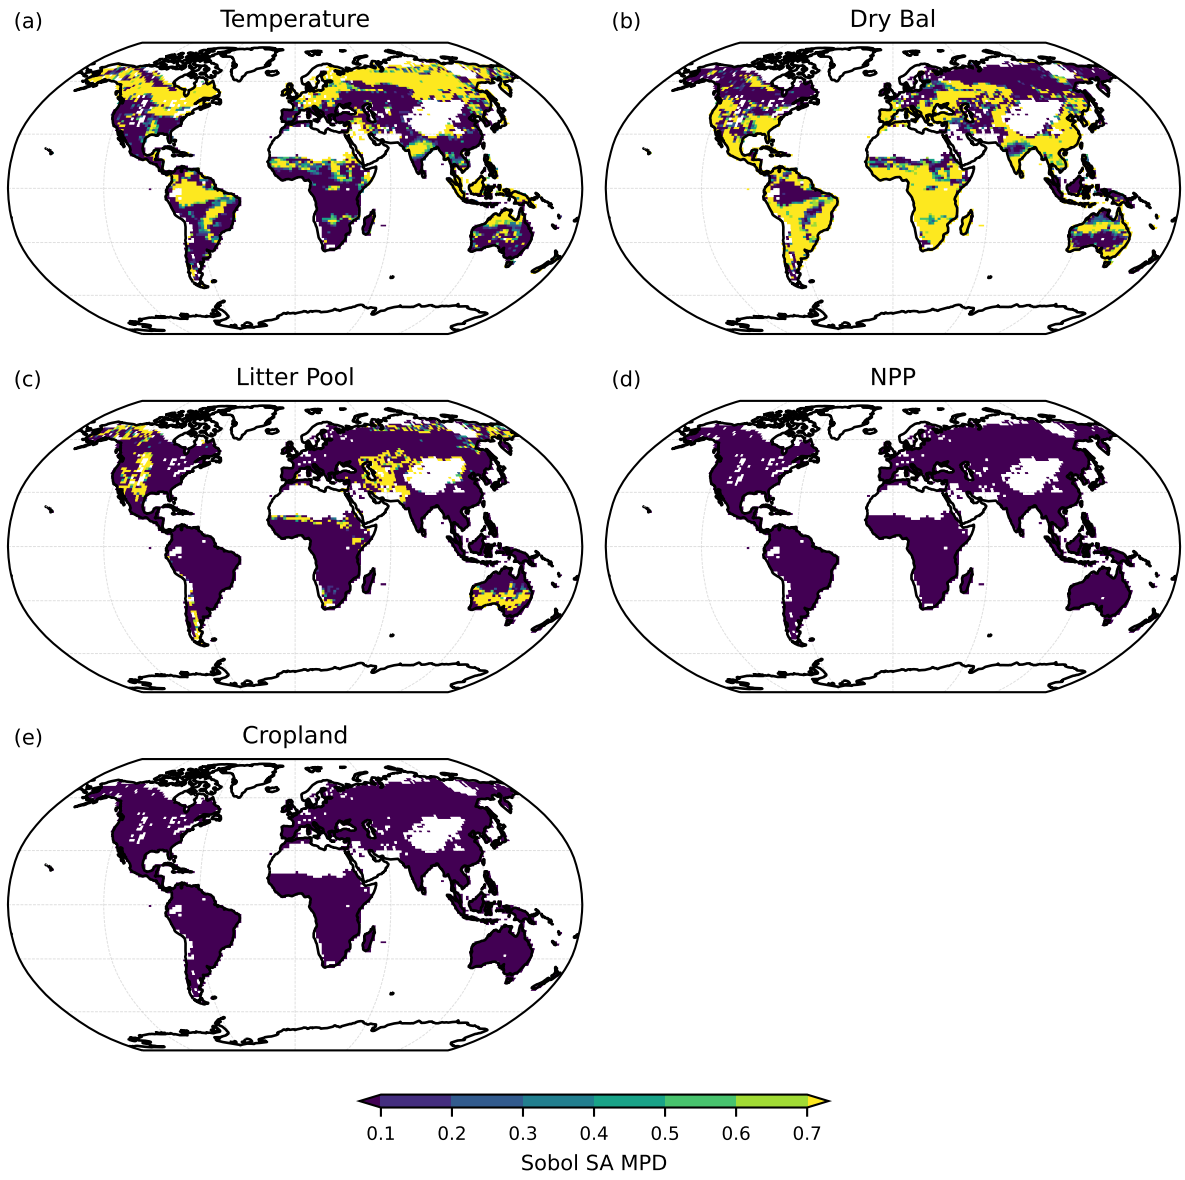


Figure 5.8: Sobol MPD S_T importances for SINFERNO-B2 data. Importances were calculated at 5342 grid cells out of 7771 land grid cells, with grid cells being omitted due to low BA. Temperature (a) and Dryness (Dry Bal) (b) are the most important globally, with the spatial distribution of the temperature significance largely contrasting that of dryness. Unlike dryness, the temperature is most important in northern regions. The spatial distribution of the significance of fuel build-up (Litter Pool) (c) largely represents the complement of the combined spatial distribution of the dryness and temperature significance. Instantaneous productivity (NPP) (d) and cropland (e) are insignificant.

The temperature sigmoid and weight parameters are globally significant for both NME and MPD metrics (Fig. 5.9a–d and 5.10a–d). The temperature weight parameter is mostly important in northern regions. On the other hand, the temperature centre parameter has a significant influence on the MPD on a global scale (Fig. 5.10b).

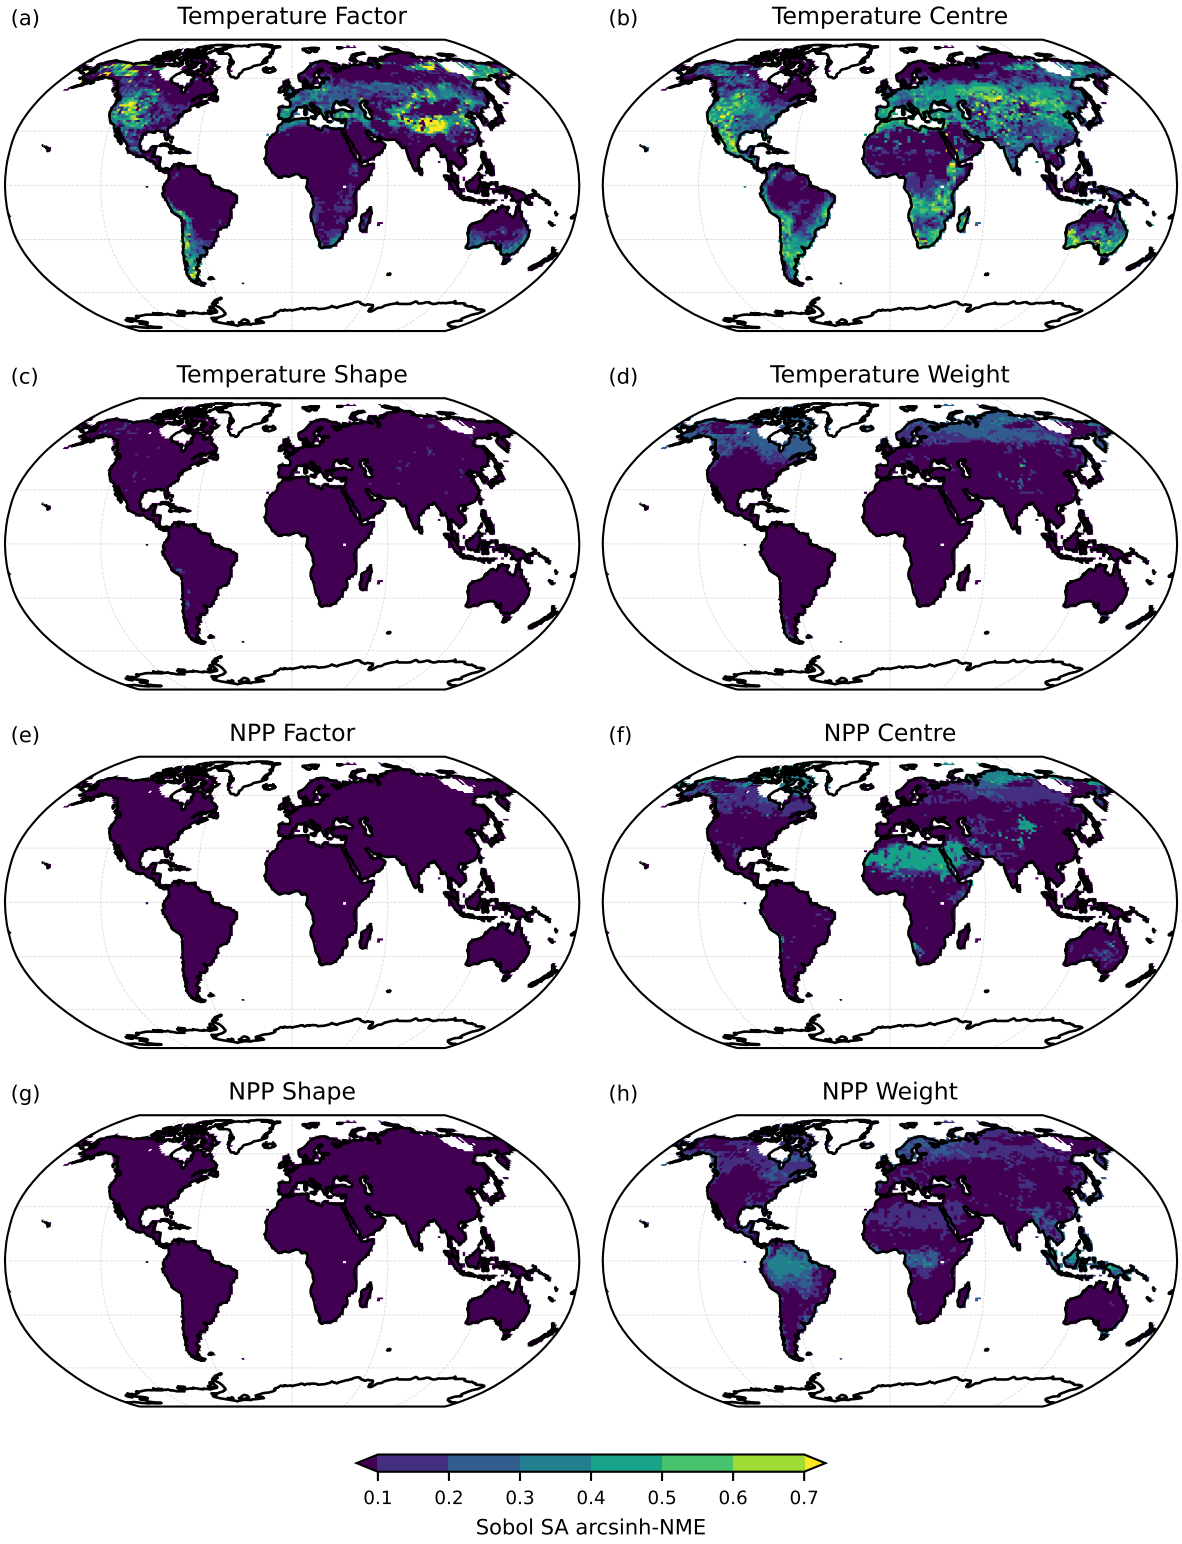


Figure 5.9: Sobol arcsinh-NME S_T importances for SINFERNO-B2 temperature (a–d) and NPP (e–h) parameters. Importances were calculated at 7243 grid cells out of 7771 land grid cells. The temperature factor (a) and centre (b) parameters are the most important, primarily in extra-tropical regions. The temperature weight (d) is exclusively important in northern, temperature-limited regions, while the NPP weight (h) parameter tends to be important in tropical NPP-limited regions. The temperature shape (c), NPP factor (e), and NPP shape (g) parameters are insignificant.

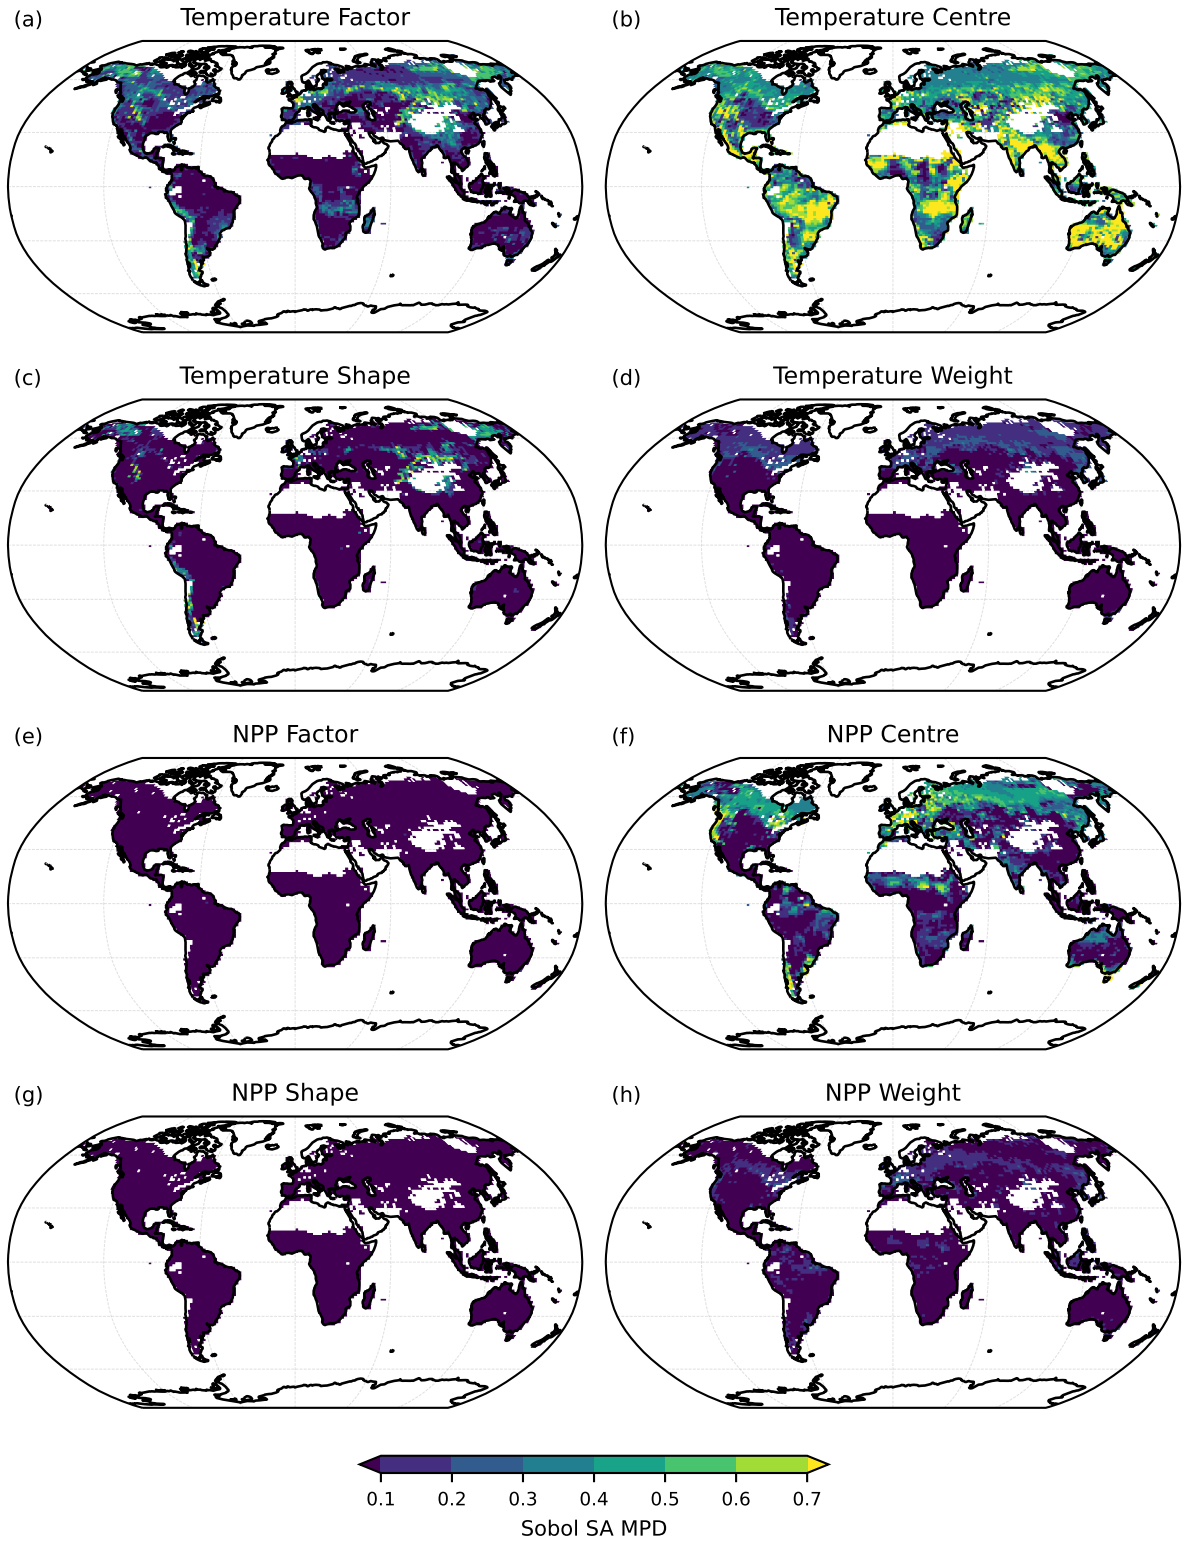


Figure 5.10: Sobol MPD S_T importances for SINFERNO-B2 temperature (a–d) and NPP parameters (e–h). Importances were calculated at 5536 grid cells out of 7771 land grid cells, with grid cells being omitted due to low BA. The temperature centre (b) parameter is the most important globally. The temperature factor (a) is mostly important in northern regions and parts of South America. The NPP centre (f) parameter is important in northern regions and the Sahel. The NPP factor (e), shape (g), and weight (h) parameters are insignificant.

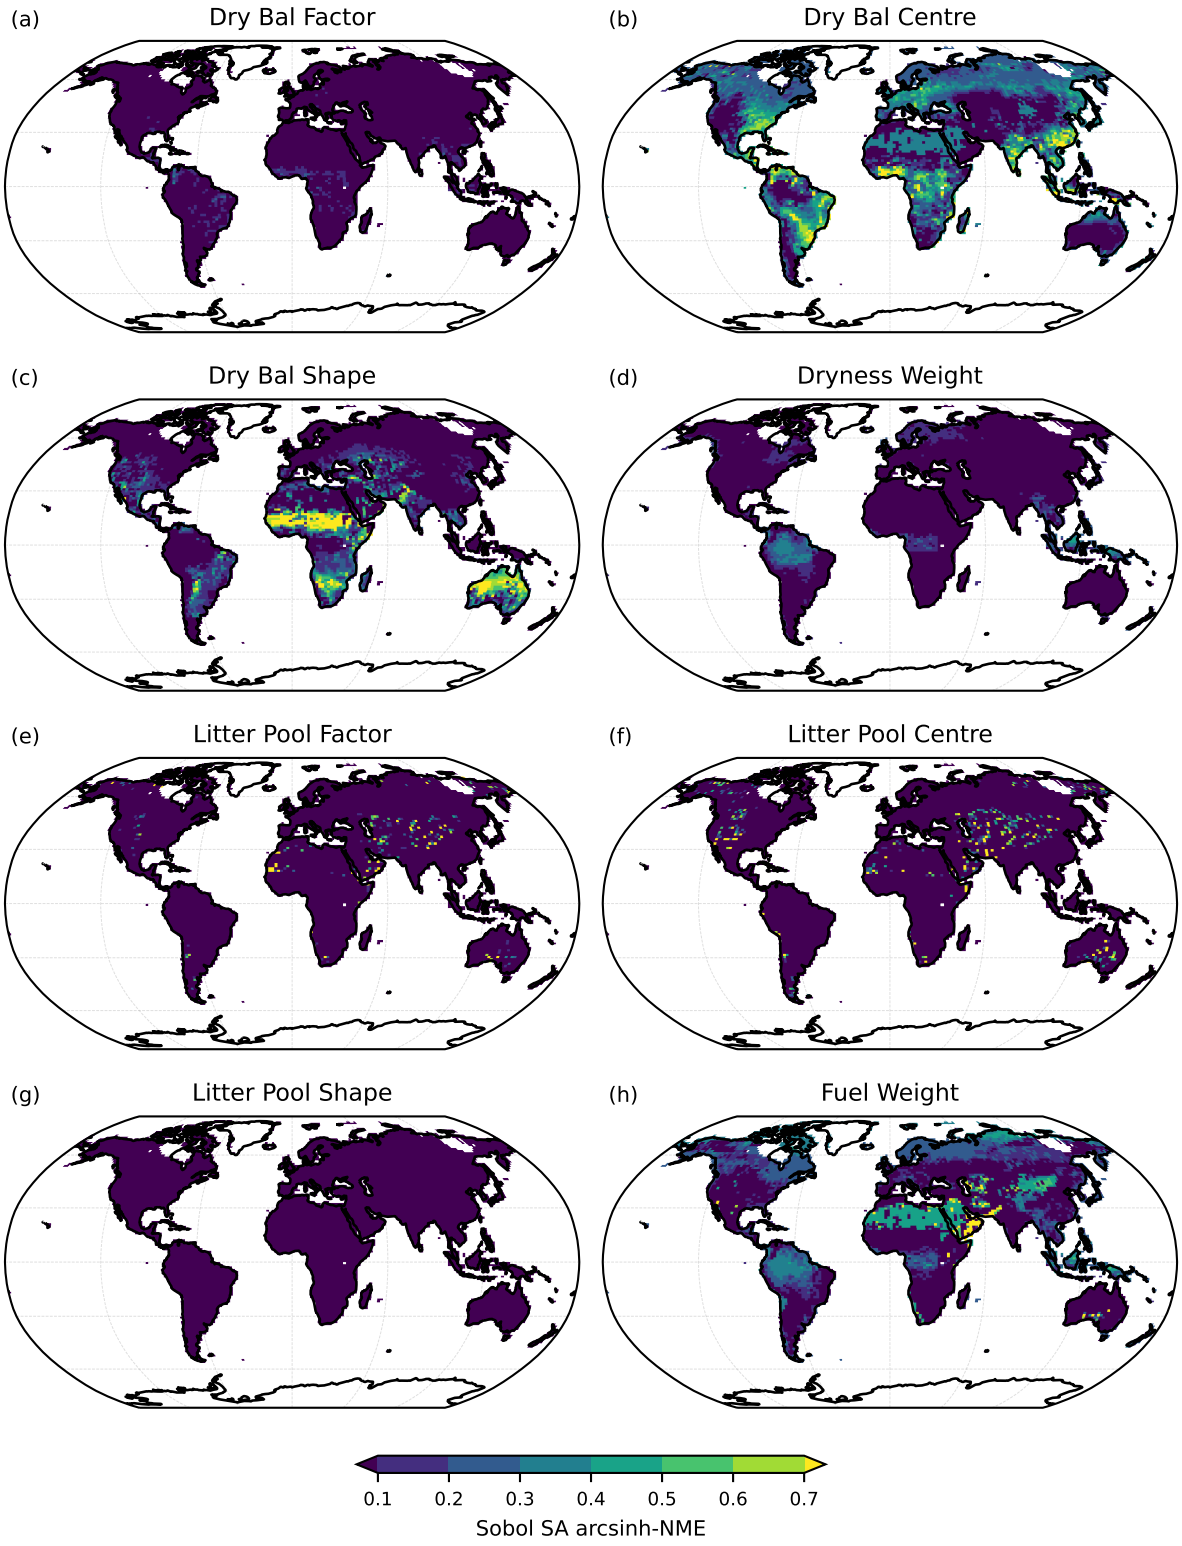


Figure 5.11: Sobol arcsinh-NME S_T importances for SINFERNO-B2 dryness (a–d) and fuel (e–h) parameters. Importances were calculated at 7243 grid cells out of 7771 land grid cells. The dryness (Dry Bal) centre (b) and shape (c) parameters are the most important globally. Although most regions are predominantly affected by the dryness centre parameter (b), certain regions, such as the Sahel, Australia, and South Africa, are more affected by the dryness shape parameter (c). All other parameters are insignificant, with the exception of the fuel weight parameter (h), which is of minor importance in northern and tropical regions.

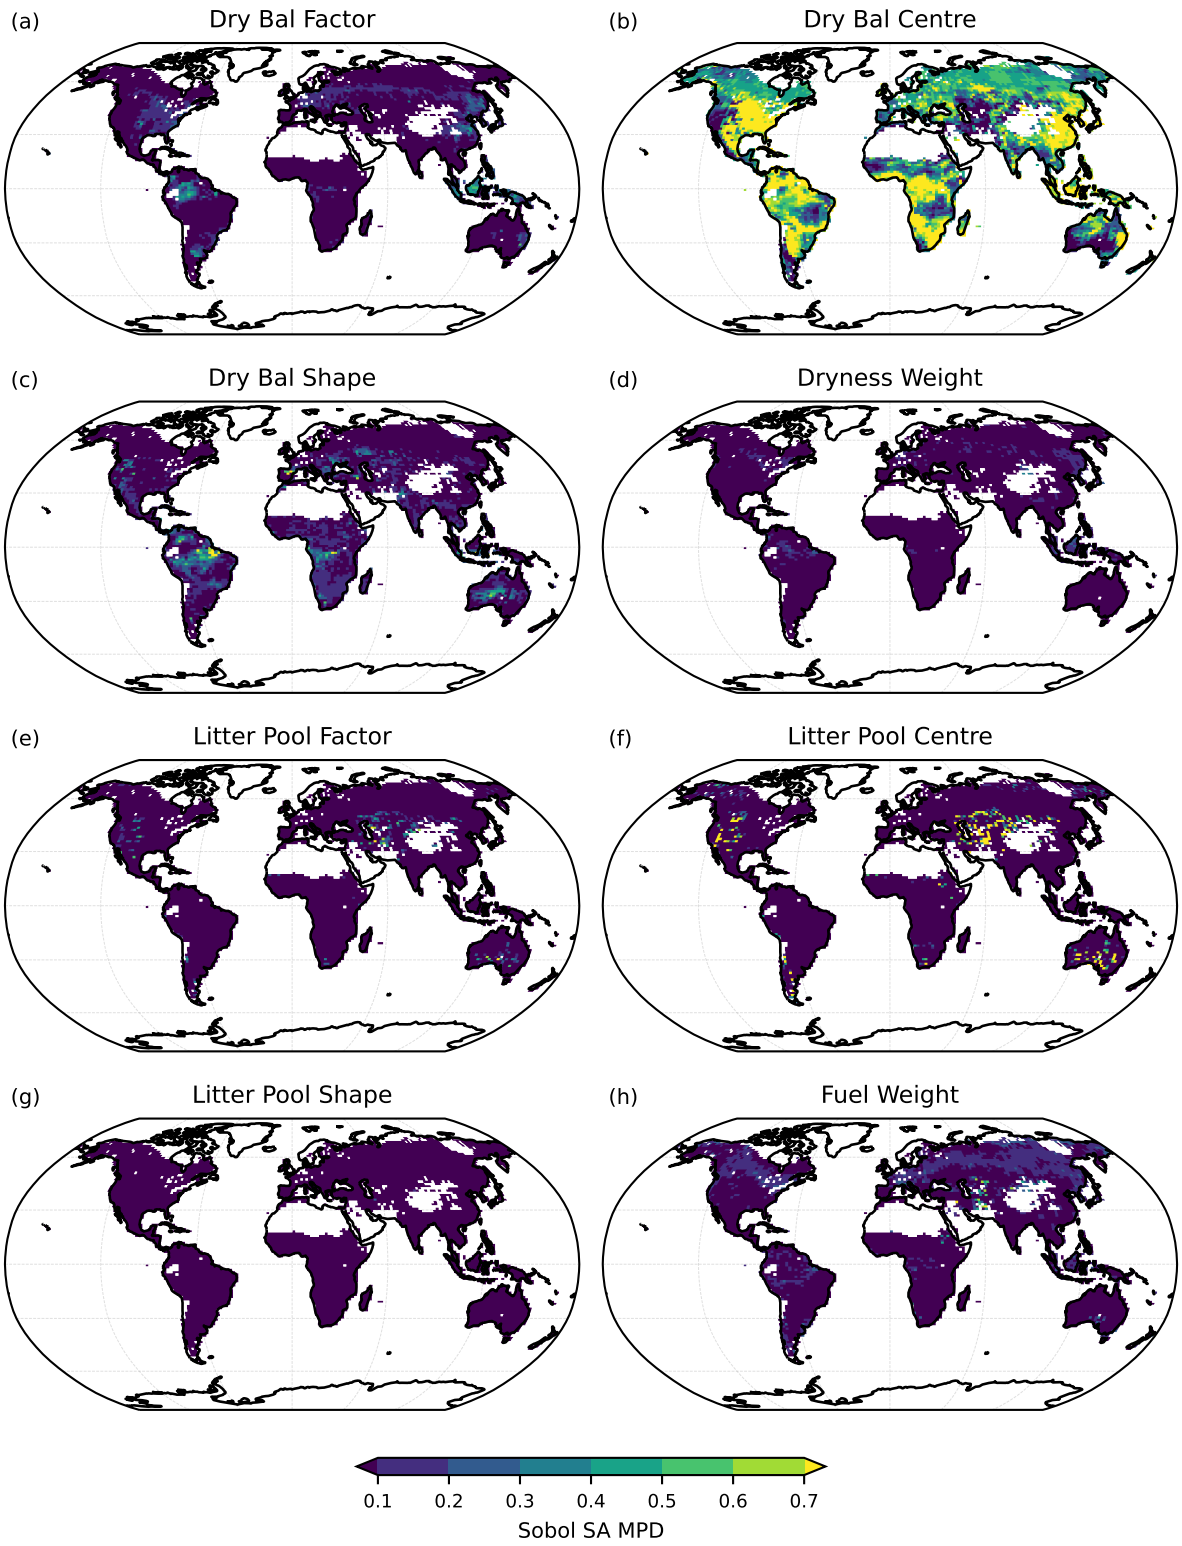


Figure 5.12: Sobol MPD S_T importances for SINFERNO-B2 dryness (a–d) and fuel (e–h) parameters. Importances were calculated at 5536 grid cells out of 7771 land grid cells, with grid cells being omitted due to low BA. The dryness (Dry Bal) centre parameter (b) is the most important globally. The dryness factor (a) and shape (c) parameters are of minor, scattered importance in some (e.g. tropical) regions. Other parameters are insignificant.

For the instantaneous productivity (Fig. 5.9e–h and 5.10e–h), the sigmoid parameters are mostly insignificant, with the exception of the centre and weight parameters for the NME and the centre parameter for the MPD. The weight parameter has a small influence in tropical regions (Fig. 5.9h). Regarding the NME, and ignoring the desert regions, which exhibit next to no BA, the sigmoid centre parameter is significant in northern regions (Fig. 5.9f), contributing to improved NME performance due to the regulation of the phase of burning. This is seen more clearly in the global influence of this parameter on the MPD (Fig. 5.10f), with the only exceptions being regions with relatively little NPP variability, such as tropical forest regions, where the modelled NPP relationship mainly tends to suppress BA. The NPP factor and shape parameters are generally unimportant.

Parameters concerning dryness, as represented by Dry Bal (Fig. 5.11a–d and 5.12a–d), are highly significant globally, with the sigmoid centre parameter being important for both the NME and MPD and the shape parameter mainly influencing the NME.

Litter pool parameters (Fig. 5.11e–h and 5.12e–h) are not significant, with the exception of the weight parameter, which predominantly has a small influence on the NME in tropical and northern regions. However, these regions do not completely overlap with regions important for fuel build-up, as shown in Fig. 5.7c and 5.8c. Although the litter pool variable was shown to be highly significant in these isolated regions, all of the litter pool parameters are much less significant in comparison.

5.3.3 Reduced complexity models

Model performance initially improves as the number of parameters increases, asymptotically approaching a minimum loss before worsening again for the most complicated models with the most parameters (see Fig. 5.13). The SINFERNO-B2 model generally performed the best across the range of the number of optimised parameters. The ideal number of parameters (see Table 5.6) was found to depend on the method employed. Using the AIC (both with and without arcsinh transformation of the data prior to the computation of the SSE) generally resulted in ideal

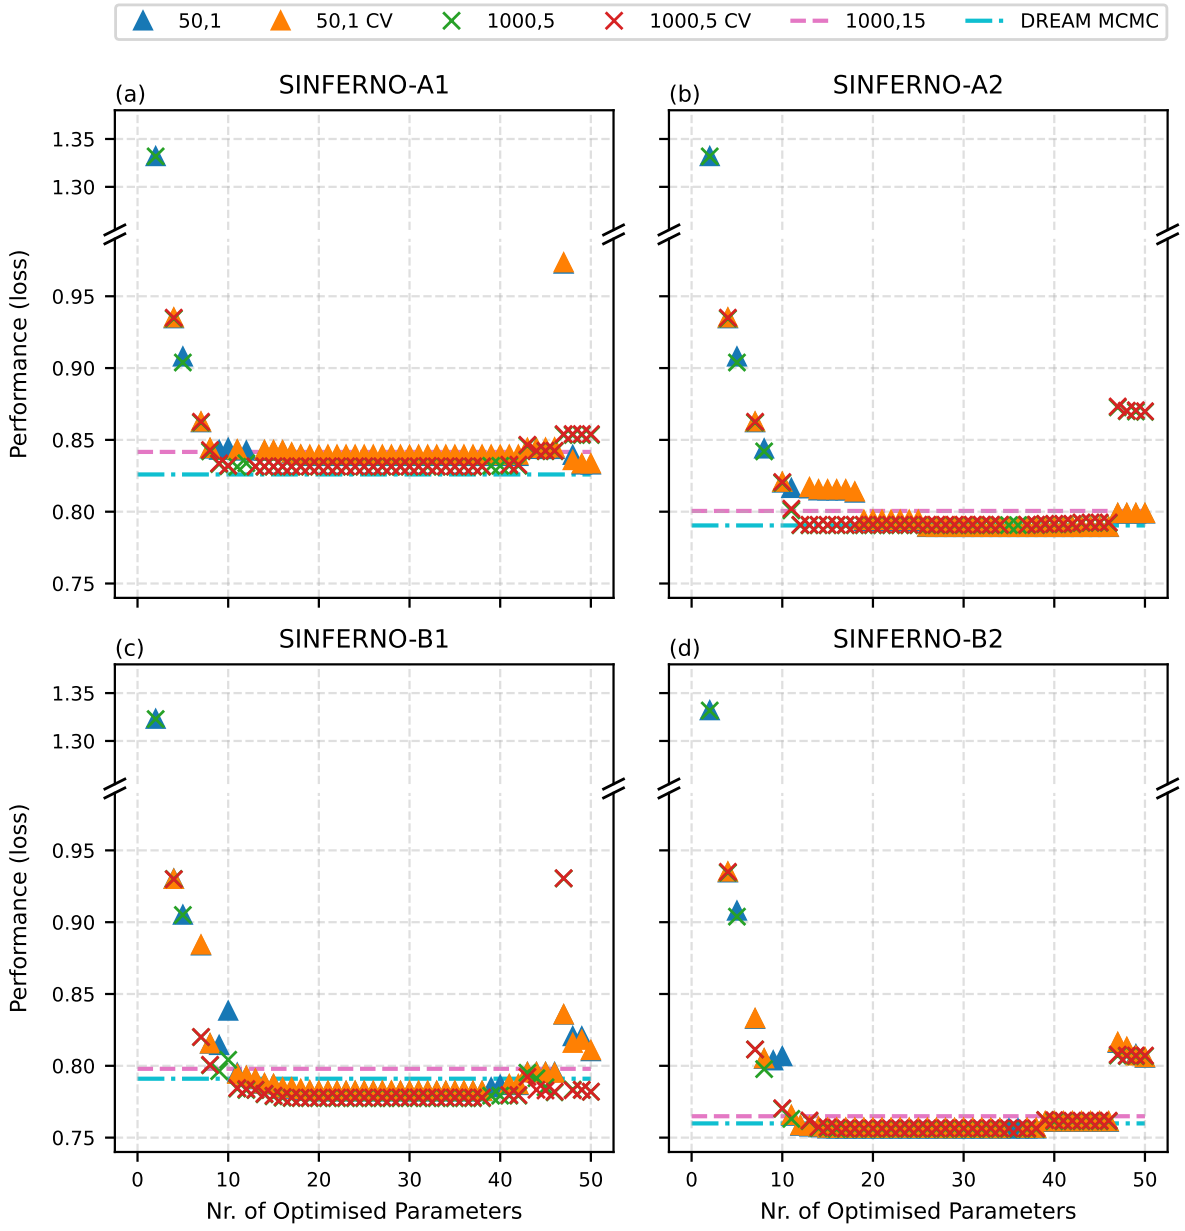


Figure 5.13: Iterative model optimisation performance, measured using the composite loss. Model performance is generally seen to improve as the number of optimised parameters increases. Relative improvement gains diminish as model complexity increases, with performance worsening as the maximum number of parameters is reached. Different optimisation parameters used: blue triangles denote a maximum of 50 local minimiser iterations with termination of the basin-hopping algorithm after 1 step yields the same minimum. Green crosses denote the higher computational cost of 1000 minimiser iterations and a minimum of 5 basin-hopping iterations. Five-fold cross-validation results are shown in orange and red. The performance of the initial basin-hopping optimisation and the DREAM MCMC optimisation is shown using dotted lines. Note the split y-axis due to the much lower performance of the models using only two parameters.

Table 5.6: Best-performing number of optimised parameters using various methods. The AIC was calculated using the SSE (both with and without prior arcsinh transformation of the data). Five-fold CV was carried out using the composite loss (arcsinh-NME + MPD). AIC analysis generally led to less complex models being preferred compared to CV. For some models, the CV analysis yielded identical results, leading to a range of ideal parameters.

Scheme	Optimisation	SSE AIC	arcsinh-SSE AIC	CV
A1	50,1	9	44	49–50
	1000,5	5	5	22
A2	50,1	7	22	40
	1000,5	5	10	15–17
B1	50,1	11	13	38
	1000,5	11	9	18
B2	50,1	8	8	28
	1000,5	11	8	17

models with relatively few parameters (∼10). On the other hand, the CV approach using the composite loss yielded very different results depending on the optimisation settings employed. For the more thorough settings of 1000 local minimiser iterations and a minimum of 5 basin-hopping iterations, models with an intermediate complexity (∼20) were selected using the CV method, while the cheaper optimisation with 50 local iterations and 1 basin-hopping iteration led to more complicated models with 28–50 parameters being preferred.

Past ∼30 parameters, increasing model complexity did not lead to large performance improvements, with performance generally worsening for models with more than ∼45 parameters. SINFERNO models optimised in a standalone fashion without knowledge of the previous minima, albeit with a higher minimum number of basin-hopping iterations (15 compared to 5; '1000,15' in Fig. 5.13), performed worse than the models trained using the iterative model optimisation at ∼30 parameters.

While Fig. 5.13 provides an overarching perspective on the influence of model complexity on SINFERNO performance, it is also instructive to look at the individual model configurations arrived at during the iterative algorithm's search (such as Fig. 5.3 and 5.4). It is consistently the case that as model complexity grows and new variables may potentially be included, dryness

and temperature are included first. Fuel is generally included next, followed, finally, by NPP. The order of inclusion in the reduced-complexity SINFERNO models provides an estimate of the relative importance of different predictors for the BA predicted by SINFERNO.

Typically, model performance improved more by including additional predictors as opposed to making the response of BA to existing predictors more PFT-group dependent (e.g. by optimising separate sigmoid parameters for each PFT-group). At around 20 optimised parameters (which represents the optimal performance according to the random CV employed), models generally included all four categories.

5.4 Discussion

5.4.1 GAM analysis

While the temperature relationships fit by the GAM model generally agreed with the underlying relationships of SINFERNO-B2, the GAM relationships decreased towards the highest temperature values for the Tree and Grass PFT-groups. This contradicts the rising or plateauing relationships fit by SINFERNO-B2 but probably reflects a lack of data towards the highest temperatures since the number of observations decreases significantly towards the extremes. Since GAMs prioritise smoothness, uncertain relationships in regions with limited data may then also unduly affect the fitted relationships in regions with sufficient data.

While the SINFERNO and GAM models only agreed on the underlying NPP relationship for shrubs, there was a partial agreement for the other PFT-groups. In regions with the highest density of NPP data (i.e. low NPP), the GAM model fit roughly constant relationships, in agreement with SINFERNO, suggesting that the flat relationships fit by the SINFERNO-B2 model primarily reflect the conditions observed for most of the data, i.e. that for most values of instantaneous Tree and Grass NPP, there is no limitation on BA.

The discrepancy between the relationships fit by the GAM and SINFERNO-B2 for the Dry Bal variable may indicate overfitting to a potential local minimum represented by the particular set of training data chosen here. This is because the relationships fit by SINFERNO-B2 between Dry

Bal and flammability appear unphysical since there is apparently no dependence of flammability on VPD and precipitation for the Tree and Shrub PFT-groups. This would be expected to result in worse performance in general since it is expected that these variables do have an influence in reality, contrary to what is suggested by the flat curves fit by SINFERNO-B2.

Since the SINFERNO and GAM models rely on different mathematical formulations with different constraints, it is not surprising to encounter only qualitative agreement between their respective relationships to flammability. However, while the GAM model often fit more oscillatory relationships than the SINFERNO-B2 model, there is an overall agreement between the GAM and SINFERNO-B2 relationships in terms of the sign and shape of the relationships. The often oscillatory nature of the GAM relationships may itself be a reflection of overfitting due to insufficient smoothing. Since the effective GAM relationships depend on all data at a given time and location due to the way the GAM was formulated, this represents another source of discrepancy between the shown GAM and SINFERNO relationships. Especially the position along the x-axis of the GAM relationships is expected to be significantly affected by the combined effect of the other variables considered since, as shown in Fig. 5.6, the value of the other variables affects the offset of the fitted relationships along the x-axis. Thus, the GAM analysis suggests that, overall, the fitted SINFERNO-B2 model is likely to be a reasonable description of the underlying processes.

5.4.2 Sobol sensitivity analysis

The litter pool input variable (which represents fuel build-up) has a particularly large regional variability, as the magnitude of the global standard deviation far exceeds the magnitude of the average total effect itself. For most variables, the significance standard deviation is at least as large as the mean global importance itself, indicating large regional variability of the Sobol indices. The large regional variability of the input data variable importance arises from the fact that burning is limited (controlled) by different factors in different biomes. For example, wet tropical biomes (such as in central Africa) are not limited by temperature but rather by

fuel drying (or lack thereof). Similarly, the phase of burning in boreal and tundra biomes is controlled by temperature, not dryness, because this is the primary limiting factor. Interestingly, the SINFERNO-B2 litter pool (fuel build-up) variable has a large effect on both the NME and MPD in northern regions such as Alaska, where the SINFERNO models perform much better than the standard INFERNO, indicating that this may be responsible for limiting the predicted magnitude of BA in these regions, resulting in a lower BA that is more in line with observations.

The overarching significance of the dryness sigmoid centre parameter supports the importance of the dry-day parameter from Chapter 3. The larger importance of the centre parameters for the MPD metric compared to the NME metric is also expected since a modification to the sigmoidal relationship centres will significantly influence the timing of peak BA, while the other sigmoid parameters are likely to have a larger influence on the BA magnitude instead.

The low importance of instantaneous productivity and cropland on both the NME and MPD may be an artefact of the method used to define the Sobol SA sampling bounds for the data variables. Since this was based on the IAV of yearly means at each grid cell, a low IAV would necessarily lead to a low significance. This is supported by Tables 5.3 and 5.4, where it is apparent that parameters related to the instantaneous NPP are significant, despite the low importance of the corresponding NPP data variable.

The high importance of the temperature weight parameter in northern regions is likely linked to the high temperature seasonality in these regions. Therefore, the response of BA to temperature is most pronounced there, leading to the elevated importance of the weight parameter. This is also linked to the relatively low importance of the temperature sigmoid parameters (factor, centre, and shape) for the NME in regions with the highest BA (e.g. the Sahel and northern Australia). In these regions, the temperature is likely generally high enough to sustain burning, even given a (reasonable) variation of the sigmoid parameters, thereby leading to the low importance of these parameters. The temperature centre parameter is significant for global MPD because changing the sigmoid centre parameter determines the time of year when the modelled relationship with temperature permits fire more so than the factor and

shape parameters, which predominantly influence the extent to which temperature changes may influence the modelled BA.

With NPP typically being high and relatively constant in tropical regions, the modelled sigmoidal relationship between NPP and BA usually acts to suppress BA. However, the effect of this suppression is modulated by the weight parameter, which, when set to 0, may turn off the suppression altogether. This may explain the influence of the NPP weight parameter as being due to the varying degree of BA suppression as the weight parameter is changed. Since the NPP factor and shape parameters are generally insignificant, this indicates that within the reasonable parameter value bounds determined using the MCMC simulation, the centre parameter is by far the most important for the relationship between NPP and BA.

Despite the fact that the factor and shape parameters control very similar aspects of the sigmoid relationships, it is interesting to note that the dryness factor is insignificant compared to the shape parameter. In addition to the model behaviour, this may be due to the Sobol analysis sampling bounds set for these variables. Mismatched importance attribution when comparing the input variables (e.g. litter pool) to their relevant parameters (SINFERNO-B2 litter pool sigmoid parameters) is likely due to the fact that choosing appropriate sampling bounds for the Sobol SA analysis is difficult. This appears to be only partially alleviated by the use of MCMC sampling to narrow down the range of parameter values supported by the observations.

The fact that interactions (as measured by the difference between the total effect, S_T, and the first-order effect, S_1) are generally small for the input data variables suggests that the effect of any given input data variable does not depend on the other input data variables. This is expected from the way the SINFERNO models are formulated since such conditional interactions (e.g. a decreasing effect of temperature for low dry-days but an increasing effect of temperature for high dry-days) cannot arise from the multiplication of several sigmoid functions. However, one special case is the fact that for certain combinations of input variables and PFT-groups (e.g. very low temperatures and the Tree PFT-group, which is associated with a weight parameter of 1.0), the SINFERNO models will never predict non-zero flammability. In this case, the effect of e.g. fuel

build-up does depend on temperature. The fact that these cases are relatively unlikely is reflected in the very low magnitude of the interactions apparent from Tables 5.1 and 5.2.

5.4.3 Reduced complexity models

The preference for additional model parameters given less local minimisation (i.e. fewer local minimisation iterations) is likely due to the effect of local minimisation on the degree of overfitting. As the number of possible minimisation iterations is increased, models tend to fit the training data better, potentially at the expense of performance on unseen data (i.e. generalisability). Similarly, an increased number of model parameters (increased number of degrees of freedom) typically allows models to fit the training data better, which may also decrease generalisability. It follows that increasing either the model complexity (number of parameters) or the minimisation effort is likely to increase model performance. At the same time, increasing either of the two too much may lead to overfitting of the training data and therefore to worse CV performance. Given this trade-off between model complexity and minimisation effort, the best CV performance, i.e. generalisability, is achieved with either a complex model but low minimisation effort or vice versa. However, as seen in Fig. 5.13, both the original and CV performance of the models remain mostly constant across the range of 20–40 parameters, indicating that only small performance differences within this range contributed to the ordering shown in Table 5.6.

Performance on the training set is expected to improve with every added parameter since the existing minimum can be further refined using the added degree(s) of freedom. However, the absence of a performance increase with increasing complexity past ∼30 parameters is likely due to the fact that the computational cost of searching for the global minimum scales very poorly with the number of dimensions (number of optimised parameters)[1]. Changing the optimisation parameters to both explore multiple minima (basin-hopping) and spend more time exploring local minima (local minimisation) (orange triangles instead of red crosses in Fig. 5.13) therefore

did not improve performance significantly.

The standalone optimisation of the SINFERNO 'full' complexity 50-parameter models performed in Chapter 4, with 1000 local minimisation iterations and 15 basin-hopping iterations, performed worse than the intermediate complexity SINFERNO models with ∼30 parameters. This was likely because the challenge of optimising the model parameters became too great for the computational time available. The available computational time was chosen heuristically to maximise performance in a reasonable time (<1 day). This illustrates the difficulty in optimising even a parsimonious and GPU-accelerated empirical fire model with ≤50 parameters. It also suggests that considerably more computing time might be required to find an improved global minimum. The improved performance of the DREAM MCMC sampler shown in Fig. 5.13 confirms this since the DREAM algorithm was able to find a better minimum than the basin-hopping algorithm, albeit using more computational resources.

The discontinuous nature of the composite loss also complicated the search for a global minimum. Depending on the (pseudo-)random parameters of the optimisation method employed (e.g the initial guesses used by the basin-hopping algorithm), the MPD error may vary abruptly as the number of valid grid cells (with non-zero predicted BA for at least one month) is affected by minute changes in model parameter values.

Both the Sobol method (see Section 5.3.2) and the predictor importances derived from the reduced complexity models using the inclusion order agree that temperature and dryness are the most important predictors. As with the SOFIA models in [Forkel et al., 2017], it was often found that including human influence (cropland in this case) did not significantly improve results. The cropland-specific parameter was typically not added until the last stages of the optimisation process, indicating its lower significance compared to the other data variables that were generally all included earlier. This is also supported by the low importance of cropland according to the

[1]This is despite the use of algorithms designed for this purpose, such as basin-hopping or genetic optimisation algorithms (see Drüke et al. [2019] for an example of fire model optimisation using a genetic algorithm). It also applies to sensitivity analysis, where a good coverage of the entire input space is crucial for accurate estimates of parameter significance [Saltelli et al., 2008]. This easily becomes prohibitive for higher dimensions, which is another motivation for grouping related variables prior to carrying out Sobol SA, as was done herein, in order to keep the number of dimensions being probed in check.

Sobol SA, as seen in Tables 5.1 and 5.2.

5.5 Conclusions

GAM analysis revealed that the relationships fit by the SINFERNO-B2 model are plausible and, therefore, that the constraints imposed upon the SINFERNO components during model optimisation are supported by the data.

Dryness was found to be the most important overall on the global climatological scale investigated, with the temperature being nearly as important. This supports the findings of Chapter 3, where the dry-day period was found to be the most important climatic driver of climatological BA. While the sigmoid centre parameters accounted for most of the significance for controlling both NME and MPD, all sigmoid parameters contributed to controlling NME, while the centre parameters were dominant in their influence on MPD. This is expected since the centre parameters determine the point in time when the modelled flammability alternates between being high or low. In addition to dryness and temperature, instantaneous NPP had a significant impact on MPD via its associated sigmoid centre parameter. Fuel build-up (via its associated weight parameter) contributed more to NME, on the other hand. Such insights into the parameters controlling different aspects of the modelled BA allow model improvement efforts to be more focused. For example, to improve phase performance, the focus should be directed towards the magnitude and seasonal cycle of instantaneous NPP and the values of the sigmoid centre parameters used throughout SINFERNO. Using map plots of the Sobol indices, regions with particularly high significance for either the phase or magnitude of the modelled BA could also be isolated and studied in greater detail.

Probing the significance of individual data variables and parameters revealed that there exists significant regional heterogeneity. For example, the temperature sigmoid weight parameter was most important for the NME and MPD in northern regions, while the temperature centre and dryness centre parameters had a near-global impact on both the NME and MPD. The dryness shape parameter had a large effect on the NME in the Sahel and Australia without a

significant impact on the MPD, while the NPP centre parameter predominantly affected the MPD performance in northern regions and the Sahel. Therefore, using the Sobol SA technique, I was able to disentangle the predictors and parameters that were most important for predicting either the magnitude or timing of BA and where they predominantly impacted predictions. However, the reliability of the Sobol significance estimates could be further improved in the future by using an improved way of setting the Sobol SA sampling bounds for data variables as opposed to the IAV of yearly mean values at each gridpoint, perhaps taking into account predictions of future variability as well.

Reduced complexity models served to illustrate the importance of including all data sources in the model for optimal performance, as well as providing a potential alternative to parameter optimisation of the full model to begin with, since model performance was improved by iteratively increasing the model complexity. Cross-validation of the models at every stage of complexity also underlined the resilience of the parsimonious SINFERNO approach to overfitting since the cross-validated performance was consistently very close to the models trained on the full dataset. However, using both cross-validation and the AIC, it was shown that the ideal model includes 20–40 parameters (compared to the full 50).

The three methods used here each provide supplementary information, which, when combined, can be used to make informed decisions about the further optimisation of an empirical model such as SINFERNO. Since GAM analysis probes the realism of the underlying relationships, i.e. sigmoidal relationships in the case of SINFERNO, it can be used to determine whether adequate representations of the underlying processes have been chosen. Both Sobol SA analysis and reduced complexity models provide a means of determining which predictor variables and model parameters are significant. Here, reduced complexity models have indicated that the ideal SINFERNO model does not include all 50 allowed parameters. Rather, a subset of $\sim$30 parameters should be chosen in order to maximise performance, ideally containing all four data variables. Sobol SA analysis has shown that most of the sigmoid shape parameters are insignificant, suggesting that these parameters may not need to be scrutinised in as much

detail as the other parameters. Therefore, a synthesis of the information provided by the three methodologies provides instructive guidance regarding which aspects of the model being investigated should be optimised.

Chapter 6

General Discussion and Conclusions

6.1 Summary

I showed in Chapter 3 that antecedent (vegetation) conditions that influence fuel build-up and drying have a significant influence on the global burnt area (BA) and therefore, by extension, flammability (see Table 6.1). The fraction of absorbed photosynthetically active radiation (FAPAR) was found to be the most significant vegetation variable compared to the leaf area index (LAI), vegetation optical depth (VOD), and solar-induced fluorescence (SIF), while the dry-day period and maximum temperature were the most important climatic variables. There are consistent differences between instantaneous and antecedent relationships with BA, such as a positive effect of antecedent FAPAR but a negative effect of instantaneous FAPAR. This is due to the influence of antecedent FAPAR on fuel build-up, with vegetation growth in the months leading up to the fire potentially increasing fuel loads, especially in biomes with a short fire return interval (FRI) and fast vegetation regrowth. Conversely, low instantaneous FAPAR suggests the prevalence of dry conditions conducive to fuel drying in the month of the fire. The relative significance of antecedent and instantaneous conditions for the prediction of BA was also quantified, with conditions on timescales $\geq$1 year generally being insignificant. For FAPAR, for example, longer timescales of around 4 months contributed more towards fuel build-up in temperate regions, while recent conditions had a larger influence on fuel drying in the tropics.

Table 6.1: BEST15 RF model variables from Chapter 3 as ranked according to the composite importance metric from Section 3.2.4. Along with their importance ranking, the variables are related to the underlying physical processes they represent. Generally, variables represent fuel build-up and drying or ignitions and/or fire suppression.

Chapter 3: BEST15 RF model		
Rank	Variable	Significance
1	FAPAR	Current-month FAPAR: fuel drying.
2	DD	Current-month dry-day period: fuel drying.
3	FAPAR 1M	One-month antecedent FAPAR: fuel build-up.
4	LAI 3M	Three-month antecedent LAI: fuel build-up.
5	CROP	Cropland cover: anthropogenic influence on ignitions and suppression.
6	MaxT	Current-month maximum temperature: fuel drying.
7	SIF 9M	Nine-month antecedent SIF: fuel build-up.
8	DD 1M	One-month antecedent dry-day period: fuel build-up and drying.
9	POPD	Population density: anthropogenic influence on ignitions and suppression.
10	DD 9M	Nine-month antecedent dry-day period: fuel build-up and drying.
11	DD 6M	Six-month antecedent dry-day period: fuel build-up and drying.
12	Lightning	Lightning: natural ignitions.
13	SIF 6M	Six-month antecedent SIF: fuel build-up.
14	DD 3M	Three-month antecedent dry-day period: fuel build-up and drying.
15	DTR	Current-month diurnal temperature range: fuel drying.

Based on these findings, a new flammability parametrisation, the Sigmoidal INFERNO (SINFERNO) model, was developed in Chapter 4. Prior to parameter optimisation, an arcsinh transformation was applied to the observed BA in order to prioritise regions with lower BA which are typically poorly represented in global fire models due to an implicit focus on regions experiencing the highest orders of magnitude of BA. The new SINFERNO models outperformed the standard INteractive Fire and Emission algoRithm for Natural envirOnments (INFERNO) in northern regions, where the phase error of the predicted BA was much reduced by the explicit consideration of new antecedent predictor variables such as a newly parametrised litter pool variable or antecedent net primary productivity (NPP). However, especially in regions with typically high BA, such as the African Sahel, the new parametrisations struggled to predict high orders of magnitude of BA. Therefore, the performance of the new parametrisations did not

Table 6.2: Top two data variables for the SINFERNO-B2 fire model according to the two Sobol SA metrics and the reduced complexity models from Chapter 5. The Dry Bal combined VPD and precipitation variable and temperature are consistently the most important.

Chapter 5: SINFERNO-B2 Fire Model Data Variables			
Rank	Sobol SA		Reduced Complexity Model
	arcsinh-NME	MPD	
1	Dry Bal	Dry Bal	Temperature
2	Temperature	Temperature	Dry Bal

improve over the standard INFERNO model everywhere. Out of the four different SINFERNO parametrisations investigated in Chapter 4, the SINFERNO-B2 parametrisation, using vapour pressure deficit (VPD) and precipitation to represent dryness, and a re-parametrised litter pool to represent fuel loads, performed the best.

The SINFERNO parametrisations were investigated in Chapter 5, with a particular focus on the best-performing SINFERNO-B2 fire model parametrisation (see Tables 6.2 and 6.3). The underlying relationships were found to be described reasonably well by the sigmoidal relationships employed for the new parametrisation. Furthermore, dryness and temperature were found to be the most important predictor variables (see Table 6.2), with instantaneous NPP contributing more to the phase performance, while fuel build-up had a greater influence on the magnitude of the predicted BA.

Analysis of the significance of the different model parameters led to the same conclusion, with dryness and temperature parameters emerging as the most important (see Table 6.3). The centre sigmoid parameters of the sigmoidal relationships that were used to link the predictor variables to BA in the SINFERNO model were found to be most significant for the phase of the predicted BA, while all four parameters contributed to the overall magnitude of the predicted BA to some degree. Since the centre parameters effectively determine the point in time at which flammability responds to the underlying predictors, they had the largest impact on the phase performance. The temperature factor parameter, which influences how strongly flammability responds to temperature, was found to be highly significant by both the Sobol and

Table 6.3: Top 10 SINFERNO-B2 fire model parameters according to the two Sobol SA metrics and the reduced complexity models from Chapter 5. According to the Sobol SA, the sigmoid centre parameters, which influence the transition point between high and low flammability, are the most important, and even more so for the phase error (MPD). The reduced complexity model analysis mostly preferred the increased optimisation of factor parameters, which influence how rapidly the flammability responds to variations in the variables. See Fig. 4.5 for an overview of the different sigmoid parameters. Both the Sobol SA and the reduced complexity modelling approaches prioritised temperature and dryness, although both Sobol SA metrics favoured dryness, while the reduced complexity modelling preferred temperature.

Chapter 5: SINFERNO-B2 Fire Model Parameters			
Rank	Sobol SA		Reduced Complexity Model
	arcsinh-NME	MPD	
1	Dry Bal Centre	Dry Bal Centre	Temperature (Factor & Centre & Shape)
2	Temperature Centre	Temperature Centre	Temperature Factor (Trees; Grass & Shrubs)
3	Fuel Weight	NPP Centre	Dryness (Factor & Centre & Shape) (Trees & Grass)
4	Temperature Factor	Temperature Factor	Dryness Factor (Trees; Grass)
5	Dry Bal Shape	Dry Bal Shape	Fuel (Factor & Centre & Shape) (Grass)
6	NPP Centre	NPP Weight	Temperature Factor (Trees & Grass; Shrubs)
7	NPP Weight	Temperature Weight	NPP (Factor & Centre & Shape) (Shrubs)
8	Temperature Weight	Fuel Weight	Fuel Centre (Grass; Shrubs)
9	Dryness Weight	Dry Bal Factor	Temperature Shape (Trees; Grass & Shrubs)
10	Dry Bal Factor	Temperature Shape	Temperature Weight (Trees & Grass)

reduced complexity modelling approaches. The sigmoid shape parameters that influence the asymmetry of the underlying relationships were generally not found to be as important as the other parameters. Even though the SINFERNO model can make use of 50 parameters to be optimised with respect to the observations, using just ∼30 variables was found to perform best, with many parameters, such as most of the sigmoid shape parameters, having been identified

as unimportant. Consequently, computational and investigative efforts can be concentrated on refining the parameters that were found to have a more significant impact.

Methodologically, Chapter 3 used accumulated local effects (ALE) plots instead of the established partial dependence plots (PDPs) in order to estimate the underlying relationships fit by the random forest (RF) machine learning models in a robust manner when given correlated input variables. Similarly, other existing techniques, such as SHapley Additive exPlanations (SHAP) values, were used to gain further insights from the empirical analyses by providing alternative means of elucidating black-box machine learning models. While earlier studies have explored fire drivers, including antecedent effects [e.g. Archibald et al., 2009, Forkel et al., 2017, Joshi and Sukumar, 2021, Haas et al., 2022], I was able to quantify the relative importance of a range of antecedent time periods on a global scale using these techniques. The importance of the different predictor variables and time periods on fuel build-up and drying could also be discerned on a regional basis, allowing explicit links to be made between controls on fire and the distribution of biomes. The composite loss and sensitivity analysis (SA) methods used in Chapters 4 and 5 also provide new ways to optimise model parameters and interrogate their significance. Using approximate Bayesian computation (ABC) alongside the Sobol SA technique allowed for more robust quantification of the parameter uncertainties supported by the data. The Sobol SA approach has previously been used to understand global models, e.g. climate models [e.g. Lee et al., 2011, Wild et al., 2020]. Here, I applied it to the empirical SINFERNO fire model in order to improve the robustness of the sensitivity analysis compared to simple local perturbations [e.g. Le Page et al., 2015, Mangeon, 2017]. For example, while the global temperature perturbations of ± 1.5 K used by Mangeon [2017] represent an analysis of the model response to a specific perturbation, Sobol SA enables an analysis of parameter significance across a wide range of input data [Saltelli et al., 2008]. Therefore, while these techniques are not new, they have been adapted in this work for the study of wildfire drivers and wildfire models and, in some cases, used jointly in order to carry out empirical analyses and empirical modelling of wildfires in greater detail with increased statistical robustness. As such, there is potential to make use of them to evaluate

future fire models as part of a more comprehensive evaluation strategy.

6.2 Future work

Chapter 3 used RF algorithms to demonstrate the importance of antecedent vegetation conditions that influence fuel build-up and antecedent conditions that influence fuel drying. However, the performance of the various RF models used could be improved by re-calibrating the model hyperparameters for each set of variables. Additionally, buffered leave-one-out cross-validation (B-LOO CV) or the combined variable importance metric developed in Chapter 3 could be used to select ideal predictor variables more robustly. The longer-term monthly analysis was also carried out using the same predictor variables that were identified as being important by the climatological analysis. Re-evaluating the predictor variables included in the longer-term monthly analysis could therefore lead to new insights that eluded the climatological analyses, such as the significance of decadal antecedent effects on fuel build-up.

Prediction of productivity and vegetation indices like FAPAR was shown to be crucial for the accurate modelling of BA. However, it was also demonstrated that FAPAR is currently poorly predicted by the Joint UK Land Environment Simulator (JULES). Empirically, I found that the seasonal cycle of these quantities is important for the prediction of fuel build-up and that these are therefore crucial variables for improving fire simulations in coupled fire models such as JULES-INFERNO. Thus, further work should be undertaken to improve the prediction of FAPAR and related quantities in land surface models (LSMs) like JULES in order to improve our fire modelling capacity.

While SINFERNO improves upon the standard INFERNO BA predictions in high latitudes, it is known that peat fires play a crucial role in these regions. In addition to improved prediction of vegetation indices like FAPAR and litterfall, which could result in improved fuel build-up representation, explicit modelling of peat and the resulting ability to directly model peat fires in JULES-INFERNO would likely further improve performance both in northern regions and in equatorial Asia. Furthermore, exploring regional- or biome-specific parametrisations instead of

using PFT-specific parameters to cater to specific biomes based on their distribution of vegetation types may yield better results depending on the selected regions and the intended scope of the resulting fire model. The robustness of the SINFERNO parametrisations could also be improved by applying the CV applied to the continuous parameters in Chapter 5 to the discrete parameters as well. This is especially relevant for the SINFERNO-B2 parametrisation since this has the most discrete parameters.

Chapters 3, 4, and 5 demonstrated the inability of simple proxies of anthropogenic influence on BA, like population density and cropland cover, to demonstrably improve BA predictions on the global climatological scales investigated. While incorporating the Human Development Index (HDI) has shown promise for improved INFERNO BA predictions [Teixeira et al., 2021], HDI (along with population density and cropland cover, amongst others) does not account for the seasonal variability of human ignition potential, and likely does not capture regional variability in human–fire interactions sufficiently. Therefore, more refined representations of human influence on fires may improve modelled BA in regions particularly affected by human intervention. For example, more agent-based [Ford et al., 2021] approaches, as suggested by Pfeiffer et al. [2013] for the SPITFIRE model, where different groups of people (such as pastoralists or farmers) exhibit different interactions with fire, may improve our representation of human–fire interactions. In the absence of such detailed treatment, the straightforward removal of grid cells with high cropland cover, similar to Knorr et al. [2016], may be preferable to an otherwise indeterminate and spatially heterogeneous relationship with measures of anthropogenic influence such as cropland cover.

While the Sobol SA used in Chapter 5 permits more robust estimates of parameter significance than local point estimates resulting from simple perturbations, knowledge of the magnitude of Sobol SA sampling bounds is still required in the same way that perturbation magnitudes are required for a simpler SA approach. Even though this is somewhat alleviated by the use of Markov chain Monte Carlo (MCMC) simulations in order to estimate sampling bounds for the parameters supported by the BA observations, estimating the sampling bounds for the underlying predictor variables proved to be challenging since using the inter-annual variability (IAV) of the

yearly mean values at each grid cell, as done in Chapter 5, invariably leads to the exaggeration of the importance of variables and locations with inherently large yearly IAV. Further work should be undertaken in order to understand the regional uncertainties inherent to the predictor variables used and how these should be integrated into a SA framework. For example, to improve the sampling bounds used for SA analysis of data parameters, one could consider global measures of variability, e.g. the standard deviation. Further, reduced-complexity models could be used to carry out MCMC and Sobol sensitivity analyses to capture how models with fewer parameters respond to variations in the input data and parameters. The techniques applied to the SINFERNO models could also be applied to the standard INFERNO model to enable further comparisons with the SINFERNO models. My results regarding the evaluation of JULES performance could also be related to an anticipated JULES-ES evaluation [Mathison et al., 2022], with e.g. overestimated tropical tree cover being one of their findings. Future improvements to JULES arising from such works, and resulting in improved modelling of vegetation distribution, should also be considered when using JULES-INFERNO to enable more accurate modelling of BA.

The techniques used in this thesis also have the potential to be applied to other fire models to aid in fire model intercomparison projects and, thus, their continued development. They include the more statistically rigorous ALE plots and the spatial autocorrelation analysis of machine learning models used in Chapter 3 and the Sobol and MCMC analyses from Chapter 5. Since these techniques rely on the analysis of large amounts of data to be statistically significant, emulators would likely be required to apply the techniques used here, Sobol SA and MCMC analysis in particular, to more complicated models like process-based fire models (and fire-enabled Earth system models (ESMs) in general), since these would otherwise take far too long to run. This has previously shown promise for climate models (see e.g. Wild et al. [2020]). Fire models could also be investigated for a small subset of grid cells (see e.g. Drüke et al. [2019] for an example of fire model optimisation in South America), potentially enabling the use of the analytical techniques employed here without the need for emulation.

6.3 Wider implications

Climate change has been shown to affect future fire regimes and vice versa, with fires having the potential to influence future climate, especially in the context of land-use change. Therefore, a greater understanding of the biophysical and anthropogenic drivers of fire and their interactions is crucial for an accurate assessment of future fire impacts on the terrestrial biosphere.

The empirical analysis of fire drivers in Chapter 3 has enabled new insights into the importance of antecedent climate and vegetation conditions on fire activity. In addition to confirming the importance of antecedent vegetation conditions for BA prediction, the use of correlation-tolerant techniques allowed the quantification of the relative importance of different antecedent time periods on both fuel build-up and drying on a global scale. By bettering our understanding of fire drivers, such insights can improve our capability to model global fire. Consequently, predictions of the impact of climate change can be refined.

Fire parametrisations that focus on regions with typically low BA, like the SINFERNO models developed in Chapter 4, are especially significant since these regions are expected to experience greater changes in BA under future climate scenarios. With the phase performance of existing fire models having been identified as an area for improvement, potentially due to an insufficient representation of vegetation impacts on flammability, the SINFERNO models that were formulated based on my empirical analyses of BA demonstrate that by utilising (antecedent) vegetation properties, performance in regions that already displayed good performance can largely be maintained while improving performance in (e.g. northern) regions that were previously poorly represented. Therefore, future fire modelling studies may be able to utilise similar insights in order to achieve comparable improvements to the predicted flammability.

Overall, the work presented here can improve the robustness of empirical fire driver analyses, individual fire model assessments, and fire model intercomparison projects, thereby contributing to the expanding field of fire science at a time when the implications of fire have come to be intensely scrutinised.

Bibliography

S. F. Abarca, K. L. Corbosiero, and T. J. Galarneau. An evaluation of the Worldwide Lightning Location Network (WWLLN) using the National Lightning Detection Network (NLDN) as ground truth. *J. Geophys. Res.*, 115(D18):D18206, Sept. 2010. ISSN 0148-0227. doi: 10.1029/2009JD013411.

J. T. Abatzoglou and C. A. Kolden. Relationships between climate and macroscale area burned in the western United States. *Int. J. Wildland Fire*, 22(7):1003–1020, Nov. 2013. ISSN 1448-5516. doi: 10.1071/WF13019.

J. T. Abatzoglou, A. P. Williams, L. Boschetti, M. Zubkova, and C. A. Kolden. Global patterns of interannual climate-fire relationships. *Glob Change Biol*, 24(11):5164–5175, Nov. 2018. ISSN 13541013. doi: 10.1111/gcb.14405.

J. T. Abatzoglou, A. P. Williams, and R. Barbero. Global Emergence of Anthropogenic Climate Change in Fire Weather Indices. *Geophys. Res. Lett.*, 46(1):326–336, Jan. 2019. ISSN 0094-8276, 1944-8007. doi: 10.1029/2018GL080959.

C. Albergel, C. Rüdiger, T. Pellarin, J.-C. Calvet, N. Fritz, F. Froissard, D. Suquia, A. Petitpa, B. Piguet, and E. Martin. From near-surface to root-zone soil moisture using an exponential filter: An assessment of the method based on in-situ observations and model simulations. *Hydrology and Earth System Sciences*, 12(6):1323–1337, Dec. 2008. ISSN 1027-5606. doi: 10.5194/hess-12-1323-2008.

A. Aldersley, S. J. Murray, and S. E. Cornell. Global and regional analysis of climate and human

drivers of wildfire. *Science of The Total Environment*, 409(18):3472–3481, Aug. 2011. ISSN 00489697. doi: 10.1016/j.scitotenv.2011.05.032.

I. Alonso-Canas and E. Chuvieco. Global burned area mapping from ENVISAT-MERIS and MODIS active fire data. *Remote Sensing of Environment*, 163:140–152, June 2015. ISSN 0034-4257. doi: 10.1016/j.rse.2015.03.011.

S. T. Alvarado, N. Andela, T. S. F. Silva, and S. Archibald. Thresholds of fire response to moisture and fuel load differ between tropical savannas and grasslands across continents. *Global Ecol Biogeogr*, 29(2):331–344, Feb. 2020. ISSN 1466-822X, 1466-8238. doi: 10.1111/geb.13034.

N. Andela, Y. Y. Liu, A. I. J. M. van Dijk, R. a. M. de Jeu, and T. R. McVicar. Global changes in dryland vegetation dynamics (1988–2008) assessed by satellite remote sensing: Comparing a new passive microwave vegetation density record with reflective greenness data. *Biogeosciences*, 10(10):6657–6676, Oct. 2013. ISSN 1726-4170. doi: 10.5194/bg-10-6657-2013.

N. Andela, D. C. Morton, L. Giglio, Y. Chen, G. R. van der Werf, P. S. Kasibhatla, R. S. DeFries, G. J. Collatz, S. Hantson, S. Kloster, D. Bachelet, M. Forrest, G. Lasslop, F. Li, S. Mangeon, J. R. Melton, C. Yue, and J. T. Randerson. A human-driven decline in global burned area. *Science*, 356(6345):1356–1362, June 2017. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aal4108.

P. L. Andrews, D. O. Loftsgaarden, and L. S. Bradshaw. Evaluation of fire danger rating indexes using logistic regression and percentile analysis. *Int. J. Wildland Fire*, 12(2):213–226, 2003. ISSN 1448-5516. doi: 10.1071/wf02059.

D. W. Apley and J. Zhu. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(4):1059–1086, June 2020. ISSN 1467-9868. doi: 10.1111/rssb.12377.

F. M. D. Araújo and L. G. Ferreira. Satellite-based automated burned area detection: A performance assessment of the MODIS MCD45A1 in the Brazilian savanna. *International*

Journal of Applied Earth Observation and Geoinformation, 36:94–102, Apr. 2015. ISSN 03032434. doi: 10.1016/j.jag.2014.10.009.

S. Archibald, D. P. Roy, B. W. van Wilgen, and R. J. Scholes. What limits fire? An examination of drivers of burnt area in Southern Africa. *Global Change Biology*, 15(3):613–630, Mar. 2009. ISSN 13541013, 13652486. doi: 10.1111/j.1365-2486.2008.01754.x.

S. Archibald, C. E. R. Lehmann, J. L. Gómez-Dans, and R. A. Bradstock. Defining pyromes and global syndromes of fire regimes. *PNAS*, 110(16):6442–6447, Apr. 2013. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1211466110.

S. Archibald, C. E. R. Lehmann, C. M. Belcher, W. J. Bond, R. A. Bradstock, A.-L. Daniau, K. G. Dexter, E. J. Forrestel, M. Greve, T. He, S. I. Higgins, W. A. Hoffmann, B. B. Lamont, D. J. McGlinn, G. R. Moncrieff, C. P. Osborne, J. G. Pausas, O. Price, B. S. Ripley, B. M. Rogers, D. W. Schwilk, M. F. Simon, M. R. Turetsky, G. R. V. der Werf, and A. E. Zanne. Biological and geophysical feedbacks with fire in the Earth system. *Environ. Res. Lett.*, 13(3): 033003, Mar. 2018. ISSN 1748-9326. doi: 10.1088/1748-9326/aa9ead.

P. Artaxo, J. V. Martins, M. A. Yamasoe, A. S. Procópio, T. M. Pauliquevis, M. O. Andreae, P. Guyon, L. V. Gatti, and A. M. C. Leal. Physical and chemical properties of aerosols in the wet and dry seasons in Rondônia, Amazonia. *Journal of Geophysical Research: Atmospheres*, 107(D20):LBA 49–1–LBA 49–14, 2002. ISSN 2156-2202. doi: 10.1029/2001JD000666.

E. C. Atwood, S. Englhart, E. Lorenz, W. Halle, W. Wiedemann, and F. Siegert. Detection and Characterization of Low Temperature Peat Fires during the 2015 Fire Catastrophe in Indonesia Using a New High-Sensitivity Fire Monitoring Satellite Sensor (FireBird). *PLOS ONE*, 11(8):e0159410, Aug. 2016. ISSN 1932-6203. doi: 10.1371/journal.pone.0159410.

V. Avitabile, M. Herold, G. B. M. Heuvelink, S. L. Lewis, O. L. Phillips, G. P. Asner, J. Armston, P. S. Ashton, L. Banin, N. Bayol, N. J. Berry, P. Boeckx, B. H. J. de Jong, B. DeVries, C. A. J. Girardin, E. Kearsley, J. A. Lindsell, G. Lopez-Gonzalez, R. Lucas, Y. Malhi, A. Morel, E. T. A. Mitchard, L. Nagy, L. Qie, M. J. Quinones, C. M. Ryan, S. J. W. Ferry, T. Sunderland, G. V.

Laurin, R. C. Gatti, R. Valentini, H. Verbeeck, A. Wijaya, and S. Willcock. An integrated pan-tropical biomass map using multiple reference datasets. *Global Change Biology*, 22(4): 1406–1420, Apr. 2016. ISSN 13541013. doi: 10.1111/gcb.13139.

R. Barbero, J. T. Abatzoglou, N. K. Larkin, C. A. Kolden, and B. Stocks. Climate change presents increased potential for very large fires in the contiguous United States. *Int. J. Wildland Fire*, 24(7):892–899, Nov. 2015. ISSN 1448-5516. doi: 10.1071/WF15083.

A. Bastarrika, E. Chuvieco, and M. P. Martín. Mapping burned areas from Landsat TM/ETM+ data with a two-phase algorithm: Balancing omission and commission errors. *Remote Sensing of Environment*, 115(4):1003–1012, Apr. 2011. ISSN 0034-4257. doi: 10.1016/j.rse.2010.12.005.

P. S. A. Beck, C. Atzberger, K. A. Høgda, B. Johansen, and A. K. Skidmore. Improved monitoring of vegetation dynamics at very high latitudes: A new method using MODIS NDVI. *Remote Sensing of Environment*, 100(3):321–334, Feb. 2006. ISSN 0034-4257. doi: 10.1016/j.rse.2005.10.021.

J. Bedia, S. Herrera, J. M. Gutiérrez, A. Benali, S. Brands, B. Mota, and J. M. Moreno. Global patterns in the sensitivity of burned area to fire-weather: Implications for climate change. *Agricultural and Forest Meteorology*, 214–215:369–379, Dec. 2015. ISSN 01681923. doi: 10.1016/j.agrformet.2015.09.002.

A. Benali, B. Mota, N. Carvalhais, D. Oom, L. M. Miller, M. L. Campagnolo, and J. M. C. Pereira. Bimodal fire regimes unveil a global-scale anthropogenic fingerprint. *Global Ecol Biogeogr*, 26(7):799–811, July 2017. ISSN 1466822X. doi: 10.1111/geb.12586.

J. Bergstra, D. Yamins, and D. Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In S. Dasgupta and D. McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 115–123, Atlanta, Georgia, USA, June 2013. PMLR.

W. C. Bessie and E. A. Johnson. The Relative Importance of Fuels and Weather on Fire Behavior in Subalpine Forests. *Ecology*, 76(3):747–762, 1995. ISSN 0012-9658. doi: 10.2307/1939341.

M. J. Best, M. Pryor, D. B. Clark, G. G. Rooney, R. L. H. Essery, C. B. Ménard, J. M. Edwards, M. A. Hendry, A. Porson, N. Gedney, L. M. Mercado, S. Sitch, E. Blyth, O. Boucher, P. M. Cox, C. S. B. Grimmond, and R. J. Harding. The Joint UK Land Environment Simulator (JULES), model description – Part 1: Energy and water fluxes. *Geoscientific Model Development*, 4(3): 677–699, Sept. 2011. ISSN 1991-959X. doi: 10.5194/gmd-4-677-2011.

I. Bistinas, S. P. Harrison, I. C. Prentice, and J. M. C. Pereira. Causal relationships versus emergent patterns in the global controls of fire frequency. *Biogeosciences*, 11(18):5087–5101, Sept. 2014. ISSN 1726-4170. doi: 10.5194/bg-11-5087-2014.

M. M. Boer, V. R. De Dios, E. Z. Stefaniak, and R. A. Bradstock. A hydroclimatic model for the distribution of fire on Earth. *Environ. Res. Commun.*, 3(3):035001, Mar. 2021. ISSN 2515-7620. doi: 10.1088/2515-7620/abec1f.

W. J. Bond and J. E. Keeley. Fire as a global 'herbivore': The ecology and evolution of flammable ecosystems. *Trends in Ecology & Evolution*, 20(7):387–394, July 2005. ISSN 0169-5347. doi: 10.1016/j.tree.2005.04.025.

W. J. Bond, G. F. Midgley, and F. I. Woodward. The importance of low atmospheric CO2 and fire in promoting the spread of grasslands and savannas. *Global Change Biology*, 9(7):973–982, 2003. ISSN 1365-2486. doi: 10.1046/j.1365-2486.2003.00577.x.

W. J. Bond, F. I. Woodward, and G. F. Midgley. The global distribution of ecosystems in a world without fire. *New Phytologist*, 165(2):525–538, 2005. ISSN 1469-8137. doi: 10.1111/j.1469-8137.2004.01252.x.

L. Boschetti, D. Roy, and C. Justice. International global burned area satellite product validation protocol part I–production and standardization of validation reference data. *Committee on Earth Observation Satellites: Maryland, MD, USA*, pages 1–11, 2009.

L. Boschetti, D. P. Roy, C. O. Justice, and M. L. Humber. MODIS–Landsat fusion for large area 30m burned area mapping. *Remote Sensing of Environment*, 161:27–42, May 2015. ISSN 0034-4257. doi: 10.1016/j.rse.2015.01.022.

L. Boschetti, S. V. Stehman, and D. P. Roy. A stratified random sampling design in space and time for regional to global scale burned area product validation. *Remote Sensing of Environment*, 186:465–478, Dec. 2016. ISSN 00344257. doi: 10.1016/j.rse.2016.09.016.

D. M. J. S. Bowman, J. K. Balch, P. Artaxo, W. J. Bond, J. M. Carlson, M. A. Cochrane, C. M. D'Antonio, R. S. DeFries, J. C. Doyle, S. P. Harrison, F. H. Johnston, J. E. Keeley, M. A. Krawchuk, C. A. Kull, J. B. Marston, M. A. Moritz, I. C. Prentice, C. I. Roos, A. C. Scott, T. W. Swetnam, G. R. van der Werf, and S. J. Pyne. Fire in the Earth System. *Science*, 324 (5926):481–484, Apr. 2009. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1163886.

D. M. J. S. Bowman, J. Balch, P. Artaxo, W. J. Bond, M. A. Cochrane, C. M. D'Antonio, R. DeFries, F. H. Johnston, J. E. Keeley, M. A. Krawchuk, C. A. Kull, M. Mack, M. A. Moritz, S. Pyne, C. I. Roos, A. C. Scott, N. S. Sodhi, and T. W. Swetnam. The human dimension of fire regimes on Earth. *Journal of Biogeography*, 38(12):2223–2236, Dec. 2011. ISSN 1365-2699. doi: 10.1111/j.1365-2699.2011.02595.x.

D. M. J. S. Bowman, G. J. Williamson, J. T. Abatzoglou, C. A. Kolden, M. A. Cochrane, and A. M. S. Smith. Human exposure and sensitivity to globally extreme wildfire events. *Nature Ecology & Evolution*, 1(3):1–6, Feb. 2017. ISSN 2397-334X. doi: 10.1038/s41559-016-0058.

L. Breiman. Random Forests. *Machine Learning*, 45(1):5–32, Oct. 2001. ISSN 1573-0565. doi: 10.1023/A:1010933404324.

L. Budic, G. Didenko, and C. F. Dormann. Squares of different sizes: Effect of geographical projection on model parameter estimates in species distribution modeling. *Ecology and Evolution*, 6(1):202–211, 2016. ISSN 2045-7758. doi: 10.1002/ece3.1838.

R. E. Bürgesser. Assessment of the World Wide Lightning Location Network (WWLLN) detection

efficiency by comparison to the Lightning Imaging Sensor (LIS): WWLLN Detection Efficiency Relative to LIS. *Q.J.R. Meteorol. Soc*, 143(708):2809–2817, Oct. 2017. ISSN 00359009. doi: 10.1002/qj.3129.

C. Burton, R. A. Betts, C. D. Jones, and K. Williams. Will Fire Danger Be Reduced by Using Solar Radiation Management to Limit Global Warming to 1.5 °C Compared to 2.0 °C? *Geophys. Res. Lett.*, 45(8):3644–3652, Apr. 2018. ISSN 00948276. doi: 10.1002/2018GL077848.

C. Burton, R. Betts, M. Cardoso, T. R. Feldpausch, A. Harper, C. D. Jones, D. I. Kelley, E. Robertson, and A. Wiltshire. Representation of fire, land-use change and vegetation dynamics in the Joint UK Land Environment Simulator vn4.9 (JULES). *Geosci. Model Dev.*, 12(1):179–193, Jan. 2019. ISSN 1991-9603. doi: 10.5194/gmd-12-179-2019.

G. Chastaing, F. Gamboa, and C. Prieur. Generalized Hoeffding-Sobol decomposition for dependent variables - application to sensitivity analysis. *Electronic Journal of Statistics*, 6 (none):2420–2448, Jan. 2012. ISSN 1935-7524, 1935-7524. doi: 10.1214/12-EJS749.

Y. Chen, D. C. Morton, N. Andela, G. R. van der Werf, L. Giglio, and J. T. Randerson. A pan-tropical cascade of fire driven by El Niño/Southern Oscillation. *Nature Clim Change*, 7 (12):906–911, Dec. 2017. ISSN 1758-678X, 1758-6798. doi: 10.1038/s41558-017-0014-8.

E. Chuvieco, M. L. Pettinari, J. Lizundia-Loiola, T. Storm, and M. Padilla Parellada. ESA Fire Climate Change Initiative (Fire_cci): MODIS Fire_cci Burned Area Pixel product, version 5.1, 2018.

E. Chuvieco, F. Mouillot, G. R. van der Werf, J. San Miguel, M. Tanase, N. Koutsias, M. García, M. Yebra, M. Padilla, I. Gitas, A. Heil, T. J. Hawbaker, and L. Giglio. Historical background and current developments for mapping burned area from satellite Earth observation. *Remote Sensing of Environment*, 225:45–64, May 2019. ISSN 00344257. doi: 10.1016/j.rse.2019.02.013.

E. Chuvieco, E. Roteta, M. Sali, D. Stroppiana, M. Boettcher, G. Kirches, T. Storm, A. Khairoun, M. L. Pettinari, M. Franquesa, and C. Albergel. Building a small fire database for Sub-Saharan

Africa from Sentinel-2 high-resolution images. *Science of The Total Environment*, 845:157139, Nov. 2022. ISSN 0048-9697. doi: 10.1016/j.scitotenv.2022.157139.

D. B. Clark, L. M. Mercado, S. Sitch, C. D. Jones, N. Gedney, M. J. Best, M. Pryor, G. G. Rooney, R. L. H. Essery, E. Blyth, O. Boucher, R. J. Harding, C. Huntingford, and P. M. Cox. The Joint UK Land Environment Simulator (JULES), model description – Part 2: Carbon fluxes and vegetation dynamics. *Geoscientific Model Development*, 4(3):701–722, Sept. 2011. ISSN 1991-959X. doi: 10.5194/gmd-4-701-2011.

Copernicus Climate Change Service (C3S). ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate . Copernicus Climate Change Service Climate Data Store (CDS). https://cds.climate.copernicus.eu/cdsapp#!/home, 2017.

C. Dankers and F. Pfisterer. *Chapter 11 PFI: Training vs. Test Data — Limitations of Interpretable Machine Learning Methods*. Department of Statistics, LMU Munich, Oct. 2020.

Dask Development Team. *Dask: Library for Dynamic Task Scheduling*, 2016.

E. A. Davidson, A. C. de Araújo, P. Artaxo, J. K. Balch, I. F. Brown, M. M. C. Bustamante, M. T. Coe, R. S. DeFries, M. Keller, M. Longo, J. W. Munger, W. Schroeder, B. S. Soares-Filho, C. M. Souza, and S. C. Wofsy. The Amazon basin in transition. *Nature*, 481(7381):321–328, Jan. 2012. ISSN 1476-4687. doi: 10.1038/nature10717.

S. H. Doerr and C. Santín. Global trends in wildfire and its impacts: Perceptions versus realities in a changing world. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1696):20150345, June 2016. doi: 10.1098/rstb.2015.0345.

C. F. Dormann, J. Elith, S. Bacher, C. Buchmann, G. Carl, G. Carré, J. R. G. Marquéz, B. Gruber, B. Lafourcade, P. J. Leitão, T. Münkemüller, C. McClean, P. E. Osborne, B. Reineking, B. Schröder, A. K. Skidmore, D. Zurell, and S. Lautenbach. Collinearity: A review of methods to deal with it and a simulation study evaluating their performance. *Ecography*, 36(1):27–46, 2013. ISSN 1600-0587. doi: 10.1111/j.1600-0587.2012.07348.x.

M. Drüke, M. Forkel, W. von Bloh, B. Sakschewski, M. Cardoso, M. Bustamante, J. Kurths, and K. Thonicke. Improving the LPJmL4-SPITFIRE vegetation–fire model for South America using satellite data. *Geoscientific Model Development*, 12(12):5029–5054, Dec. 2019. ISSN 1991-959X. doi: 10.5194/gmd-12-5029-2019.

T. J. Duff and T. D. Penman. Determining the likelihood of asset destruction during wildfires: Modelling house destruction with fire simulator outputs and local-scale landscape properties. *Safety Science*, 139:105196, July 2021. ISSN 0925-7535. doi: 10.1016/j.ssci.2021.105196.

V. Fehlenberg, M. Baumann, N. Gasparri, M. Piquer-Rodriguez, G. Gavier-Pizarro, and T. Kuemmerle. The role of soybean production as an underlying driver of deforestation in the South American Chaco. *Global Environmental Change*, 45:24–34, 2017. ISSN 0959-3780. doi: 10.1016/j.gloenvcha.2017.05.001.

A. E. S. Ford, S. P. Harrison, Y. Kountouris, J. D. A. Millington, J. Mistry, O. Perkins, S. S. Rabin, G. Rein, K. Schreckenberg, C. Smith, T. E. L. Smith, and K. Yadav. Modelling Human-Fire Interactions: Combining Alternative Perspectives and Approaches. *Frontiers in Environmental Science*, 9:418, 2021. ISSN 2296-665X. doi: 10.3389/fenvs.2021.649835.

M. Forkel, W. Dorigo, G. Lasslop, I. Teubner, E. Chuvieco, and K. Thonicke. A data-driven approach to identify controls on global fire activity from satellite and climate observations (SOFIA V1). *Geoscientific Model Development*, 10(12):4443–4476, Dec. 2017. ISSN 1991-9603. doi: 10.5194/gmd-10-4443-2017.

M. Forkel, N. Andela, S. P. Harrison, G. Lasslop, M. van Marle, E. Chuvieco, W. Dorigo, M. Forrest, S. Hantson, A. Heil, F. Li, J. Melton, S. Sitch, C. Yue, and A. Arneth. Emergent relationships with respect to burned area in global satellite observations and fire-enabled vegetation models. *Biogeosciences*, 16(1):57–76, Jan. 2019a. ISSN 1726-4170. doi: 10.5194/bg-16-57-2019.

M. Forkel, W. Dorigo, G. Lasslop, E. Chuvieco, S. Hantson, A. Heil, I. Teubner, K. Thonicke, and S. P. Harrison. Recent global and regional trends in burned area and their compensating

environmental controls. *Environ. Res. Commun.*, 1(5):051005, June 2019b. ISSN 2515-7620. doi: 10.1088/2515-7620/ab25d2.

E. W. Fox, R. A. Hill, S. G. Leibowitz, A. R. Olsen, D. J. Thornbrugh, and M. H. Weber. Assessing the accuracy and stability of variable selection methods for random forest modeling in ecology. *Environ Monit Assess*, 189(7):316, July 2017. ISSN 0167-6369. doi: 10.1007/s10661-017-6025-0.

J. H. Friedman. Greedy function approximation: A gradient boosting machine. *Ann. Statist.*, 29 (5):1189–1232, Oct. 2001. ISSN 0090-5364. doi: 10.1214/aos/1013203451.

Z. Fu, P. Ciais, I. C. Prentice, P. Gentine, D. Makowski, A. Bastos, X. Luo, J. K. Green, P. C. Stoy, H. Yang, and T. Hajima. Atmospheric dryness reduces photosynthesis along a large range of soil water deficits. *Nat Commun*, 13(1):989, Feb. 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-28652-7.

Q. Ge and M. Menendez. Exploring the variance contributions of correlated model parameters: A sampling-based approach and its application in traffic simulation models. *Applied Mathematical Modelling*, 97:438–462, Sept. 2021. ISSN 0307-904X. doi: 10.1016/j.apm.2021.04.012.

H. K. Gibbs, L. Rausch, J. Munger, I. Schelly, D. C. Morton, P. Noojipady, B. Soares-Filho, P. Barreto, L. Micol, and N. F. Walker. Brazil's Soy Moratorium. *Science*, 347(6220):377–378, Jan. 2015. doi: 10.1126/science.aaa0181.

L. Giglio and D. P. Roy. On the outstanding need for a long-term, multi-decadal, validated and quality assessed record of global burned area: Caution in the use of Advanced Very High Resolution Radiometer data. *Science of Remote Sensing*, 2:100007, Dec. 2020. ISSN 2666-0172. doi: 10.1016/j.srs.2020.100007.

L. Giglio, J. T. Randerson, and G. R. van der Werf. Analysis of daily, monthly, and annual burned area using the fourth-generation global fire emissions database (GFED4). *Journal of Geophysical Research: Biogeosciences*, 118(1):317–328, 2013. ISSN 2169-8961. doi: 10.1002/jgrg.20042.

L. Giglio, L. Boschetti, D. P. Roy, M. L. Humber, and C. O. Justice. The Collection 6 MODIS burned area mapping algorithm and product. *Remote Sensing of Environment*, 217:72–85, Nov. 2018. ISSN 00344257. doi: 10.1016/j.rse.2018.08.005.

A. Goldstein, A. Kapelner, J. Bleich, and E. Pitkin. Peeking Inside the Black Box: Visualizing Statistical Learning With Plots of Individual Conditional Expectation. *Journal of Computational and Graphical Statistics*, 24(1):44–65, Jan. 2015. ISSN 1061-8600. doi: 10.1080/10618600.2014.907095.

J. Gomez-Dans, A. Spessa, M. Wooster, and P. Lewis. A Sensitivity Analysis Study of the SPITFIRE Fire Model. Preprint, Life Sciences, Oct. 2018.

M. Goss, D. L. Swain, J. T. Abatzoglou, A. Sarhadi, C. Kolden, A. P. Williams, and N. S. Diffenbaugh. Climate change is increasing the risk of extreme autumn wildfire conditions across California. *Environ. Res. Lett.*, 15(9):094016, Aug. 2020. ISSN 1748-9326. doi: 10.1088/1748-9326/ab83a7.

B. P. Graves, T. J. Ralph, P. P. Hesse, K. E. Westaway, T. Kobayashi, P. S. Gadd, and D. Mazumder. Macro-charcoal accumulation in floodplain wetlands: Problems and prospects for reconstruction of fire regimes and environmental conditions. *PLOS ONE*, 14(10):e0224011, Oct. 2019. ISSN 1932-6203. doi: 10.1371/journal.pone.0224011.

J.-M. Grégoire, H. D. Eva, A. S. Belward, I. Palumbo, D. Simonetti, A. Brink, J.-M. Grégoire, H. D. Eva, A. S. Belward, I. Palumbo, D. Simonetti, and A. Brink. Effect of land-cover change on Africa's burnt area. *Int. J. Wildland Fire*, 22(2):107–120, Sept. 2012. ISSN 1448-5516, 1448-5516. doi: 10.1071/WF11142.

G. Griffin, N. Price, and H. Portlock. Wildfires in the central Australian rangelands, 1970-1980. *Journal of Environmental Management*, 17(4):311–323, 1983.

O. Haas, I. C. Prentice, and S. P. Harrison. Global environmental controls on wildfire burnt

area, size, and intensity. *Environ. Res. Lett.*, 17(6):065004, May 2022. ISSN 1748-9326. doi: 10.1088/1748-9326/ac6a69.

S. Hantson, M. Padilla, D. Corti, and E. Chuvieco. Strengths and weaknesses of MODIS hotspots to characterize global fire occurrence. *Remote Sensing of Environment*, 131:152–159, Apr. 2013. ISSN 00344257. doi: 10.1016/j.rse.2012.12.004.

S. Hantson, A. Arneth, S. P. Harrison, D. I. Kelley, I. C. Prentice, S. S. Rabin, S. Archibald, F. Mouillot, S. R. Arnold, P. Artaxo, D. Bachelet, P. Ciais, M. Forrest, P. Friedlingstein, T. Hickler, J. O. Kaplan, S. Kloster, W. Knorr, G. Lasslop, F. Li, S. Mangeon, J. R. Melton, A. Meyn, S. Sitch, A. Spessa, G. R. van der Werf, A. Voulgarakis, and C. Yue. The status and challenge of global fire modelling. *Biogeosciences*, 13(11):3359–3375, June 2016. ISSN 1726-4189. doi: 10.5194/bg-13-3359-2016.

S. Hantson, D. I. Kelley, A. Arneth, S. P. Harrison, S. Archibald, D. Bachelet, M. Forrest, T. Hickler, G. Lasslop, F. Li, S. Mangeon, J. R. Melton, L. Nieradzik, S. S. Rabin, I. C. Prentice, T. Sheehan, S. Sitch, L. Teckentrup, A. Voulgarakis, and C. Yue. Quantitative assessment of fire and vegetation properties in historical simulations with fire-enabled vegetation models from the Fire Model Intercomparison Project. Preprint, Biogeosciences, Jan. 2020.

I. Harris, P. Jones, T. Osborn, and D. Lister. Updated high-resolution grids of monthly climatic observations – the CRU TS3.10 Dataset. *Int. J. Climatol.*, 34(3):623–642, Mar. 2014. ISSN 08998418. doi: 10.1002/joc.3711.

S. P. Harrison, I. C. Prentice, K. J. Bloomfield, N. Dong, M. Forkel, M. Forrest, R. K. Ningthoujam, A. Pellegrini, Y. Shen, M. Baudena, A. W. Cardoso, J. C. Huss, J. Joshi, I. Oliveras, J. G. Pausas, and K. J. Simpson. Understanding and modelling wildfire regimes: An ecological perspective. *Environ. Res. Lett.*, 16(12):125008, Dec. 2021. ISSN 1748-9326. doi: 10.1088/1748-9326/ac39be.

S. P. Harrison, R. Villegas-Diaz, E. Cruz-Silva, D. Gallagher, D. Kesner, P. Lincoln, Y. Shen, L. Sweeney, D. Colombaroli, A. Ali, C. Barhoumi, Y. Bergeron, T. Blyakharchuk, P. Bobek,

R. Bradshaw, J. L. Clear, S. Czerwiński, A.-L. Daniau, J. Dodson, K. J. Edwards, M. E. Edwards, A. Feurdean, D. Foster, K. Gajewski, M. Gałka, M. Garneau, T. Giesecke, G. Gil Romera, M. P. Girardin, D. Hoefer, K. Huang, J. Inoue, E. Jamrichová, N. Jasiunas, W. Jiang, G. Jiménez-Moreno, M. Karpińska-Kołaczek, P. Kołaczek, N. Kuosmanen, M. Lamentowicz, M. Lavoie, F. Li, J. Li, O. Lisitsyna, J. A. López-Sáez, R. Luelmo-Lautenschlaeger, G. Magnan, E. K. Magyari, A. Maksims, K. Marcisz, E. Marinova, J. Marlon, S. Mensing, J. Miroslaw-Grabowska, W. Oswald, S. Pérez-Díaz, R. Pérez-Obiol, S. Piilo, A. Poska, X. Qin, C. C. Remy, P. J. H. Richard, S. Salonen, N. Sasaki, H. Schneider, W. Shotyk, M. Stancikaite, D. Šteinberga, N. Stivrins, H. Takahara, Z. Tan, L. Trasune, C. E. Umbanhowar, M. Väliranta, J. Vassiljev, X. Xiao, Q. Xu, X. Xu, E. Zawisza, Y. Zhao, Z. Zhou, and J. Paillard. The Reading Palaeofire Database: An expanded global resource to document changes in fire regimes from sedimentary charcoal records. *Earth System Science Data*, 14(3):1109–1124, Mar. 2022. ISSN 1866-3508. doi: 10.5194/essd-14-1109-2022.

T. Hastie and R. Tibshirani. *Generalized Additive Models*. Chapman & Hall/CRC, Boca Raton, Fla, 1999. ISBN 978-0-412-34390-2.

T. J. Hawbaker, M. K. Vanderhoof, Y.-J. Beal, J. D. Takacs, G. L. Schmidt, J. T. Falgout, B. Williams, N. M. Fairaux, M. K. Caldwell, J. J. Picotte, S. M. Howard, S. Stitt, and J. L. Dwyer. Mapping burned areas using dense time-series of Landsat data. *Remote Sensing of Environment*, 198:504–522, Sept. 2017. ISSN 0034-4257. doi: 10.1016/j.rse.2017.06.027.

G. P. Hempson, C. L. Parr, S. Archibald, T. M. Anderson, C. J. C. Mustaphi, A. P. Dobson, J. E. Donaldson, T. A. Morrison, J. Probert, and C. M. Beale. Continent-level drivers of African pyrodiversity. *Ecography*, 41(6):889–899, June 2018. ISSN 0906-7590, 1600-0587. doi: 10.1111/ecog.03109.

T. Hengl. Global Maps of Potential Natural Vegetation at 1 km resolution, 2018.

T. Hengl, M. G. Walsh, J. Sanderman, I. Wheeler, S. P. Harrison, and I. C. Prentice. Global

mapping of potential natural vegetation: An assessment of machine learning algorithms for estimating land potential. *PeerJ*, 6:e5457, Aug. 2018. ISSN 2167-8359. doi: 10.7717/peerj.5457.

J. Herman and W. Usher. SALib: An open-source Python library for Sensitivity Analysis. *Journal of Open Source Software*, 2(9):97, Jan. 2017. ISSN 2475-9066. doi: 10.21105/joss.00097.

S. C. Herring, M. P. Hoerling, J. P. Kossin, T. C. Peterson, and P. A. Stott. Explaining Extreme Events of 2014 from a Climate Perspective. *Bull. Amer. Meteor. Soc.*, 96(12):S1–S172, Dec. 2015. ISSN 0003-0007. doi: 10.1175/BAMS-ExplainingExtremeEvents2014.1.

H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. Rosnay, I. Rozum, F. Vamborg, S. Villaume, and J.-N. Thépaut. The ERA5 global reanalysis. *Q.J.R. Meteorol. Soc.*, 146(730):1999–2049, July 2020. ISSN 0035-9009, 1477-870X. doi: 10.1002/qj.3803.

P. E. Higuera, J. T. Abatzoglou, J. S. Littell, and P. Morgan. The Changing Strength and Nature of Fire-Climate Relationships in the Northern Rocky Mountains, U.S.A., 1902-2008. *PLOS ONE*, 10(6):e0127563, June 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0127563.

T. Homma and A. Saltelli. Importance measures in global sensitivity analysis of nonlinear models. *Reliability Engineering & System Safety*, 52(1):1–17, Apr. 1996. ISSN 0951-8320. doi: 10.1016/0951-8320(96)00002-6.

C. Hong, J. A. Burney, J. Pongratz, J. E. M. S. Nabel, N. D. Mueller, R. B. Jackson, and S. J. Davis. Global and regional drivers of land-use emissions in 1961–2017. *Nature*, 589(7843): 554–561, Jan. 2021. ISSN 1476-4687. doi: 10.1038/s41586-020-03138-y.

G. Hooker. Generalized Functional ANOVA Diagnostics for High-Dimensional Functions of

Dependent Variables. *Journal of Computational and Graphical Statistics*, 16(3):709–732, Sept. 2007. ISSN 1061-8600. doi: 10.1198/106186007X237892.

G. Hooker and L. Mentch. Please Stop Permuting Features: An Explanation and Alternatives. *arXiv:1905.03151 [cs, stat]*, May 2019.

R. A. Houghton, D. L. Skole, C. A. Nobre, J. L. Hackler, K. T. Lawrence, and W. H. Chomentowski. Annual fluxes of carbon from deforestation and regrowth in the Brazilian Amazon. *Nature*, 403(6767):301–304, Jan. 2000. ISSN 1476-4687. doi: 10.1038/35002062.

T. Houska, P. Kraft, A. Chamorro-Chavez, and L. Breuer. SPOTting Model Parameters Using a Ready-Made Python Package. *PLOS ONE*, 10(12):e0145180, Dec. 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0145180.

T. Houska, Ben, P. Kraft, H. Agrawal, F. U. Jehn, P. Lauer, Iacopo, A. Vinayak, kbstn, gitporst, J. Castilla-Rho, J-M, Cecile, C. Krause, J. Lee, T. Hallouin, and Y. Wu. thouska/spotpy: v.1.5.16. Zenodo. https://doi.org/10.5281/zenodo.6676079, June 2022.

V. Huijnen, M. J. Wooster, J. W. Kaiser, D. L. A. Gaveau, J. Flemming, M. Parrington, A. Inness, D. Murdiyarso, B. Main, and M. van Weele. Fire carbon emissions over maritime southeast Asia in 2015 largest since 1997. *Sci Rep*, 6(1):26886, May 2016. ISSN 2045-2322. doi: 10.1038/srep26886.

M. L. Humber, L. Boschetti, L. Giglio, and C. O. Justice. Spatial and temporal intercomparison of four global burned area products. *International Journal of Digital Earth*, 12(4):460–484, Apr. 2019. ISSN 1753-8947, 1753-8955. doi: 10.1080/17538947.2018.1433727.

J. D. Hunter. Matplotlib: A 2D graphics environment. *Computing in science & engineering*, 9 (3):90–95, 2007.

V. Iglesias, G. I. Yospin, and C. Whitlock. Reconstruction of fire regimes through integrated paleoecological proxy data and ecological modeling. *Front. Plant Sci.*, 5, 2015. ISSN 1664-462X. doi: 10.3389/fpls.2014.00785.

T. Iwanaga, W. Usher, and J. Herman. Toward SALib 2.0: Advancing the accessibility and interpretability of global sensitivity analyses. *Socio-Environmental Systems Modelling*, 4:18155, May 2022. doi: 10.18174/sesmo.18155.

J. Jacques, C. Lavergne, and N. Devictor. Sensitivity analysis in presence of model uncertainty and correlated inputs. *Reliability Engineering & System Safety*, 91(10):1126–1134, Oct. 2006. ISSN 0951-8320. doi: 10.1016/j.ress.2005.11.047.

M. E. Jenkins, M. Bedward, O. Price, and R. A. Bradstock. Modelling Bushfire Fuel Hazard Using Biophysical Parameters. *Forests*, 11(9):925, Sept. 2020. doi: 10.3390/f11090925.

F. Johnston and D. Bowman. Bushfire Smoke: An Exemplar of Coupled Human and Natural Systems. *Geographical Research*, 52(1):45–54, 2014. ISSN 1745-5871. doi: 10.1111/1745-5871.12028.

F. H. Johnston. Understanding and managing the health impacts of poor air quality from landscape fires. *Medical Journal of Australia*, 207(6):229–230, 2017. ISSN 1326-5377. doi: 10.5694/mja17.00072.

F. H. Johnston, S. B. Henderson, Y. Chen, J. T. Randerson, M. Marlier, R. S. DeFries, P. Kinney, D. M. Bowman, and M. Brauer. Estimated Global Mortality Attributable to Smoke from Landscape Fires. *Environ Health Perspect*, 120(5):695–701, May 2012. ISSN 0091-6765. doi: 10.1289/ehp.1104422.

W. M. Jolly, M. A. Cochrane, P. H. Freeborn, Z. A. Holden, T. J. Brown, G. J. Williamson, and D. M. J. S. Bowman. Climate-induced variations in global wildfire danger from 1979 to 2013. *Nature Communications*, 6(1):7537, July 2015. ISSN 2041-1723. doi: 10.1038/ncomms8537.

M. W. Jones, J. T. Abatzoglou, S. Veraverbeke, N. Andela, G. Lasslop, M. Forkel, A. J. P. Smith, C. Burton, R. A. Betts, G. R. van der Werf, S. Sitch, J. G. Canadell, C. Santín, C. Kolden, S. H. Doerr, and C. Le Quéré. Global and Regional Trends and Drivers of Fire

Under Climate Change. *Reviews of Geophysics*, 60(3):e2020RG000726, 2022. ISSN 1944-9208. doi: 10.1029/2020RG000726.

J. Joshi and R. Sukumar. Improving prediction and assessment of global fires using multilayer neural networks. *Scientific Reports*, 11(1):3295, Feb. 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-81233-4.

M. N. Kahiu and N. P. Hanan. Fire in sub-Saharan Africa: The fuel, cure and connectivity hypothesis. *Global Ecology and Biogeography*, 27(8):946–957, 2018. ISSN 1466-8238. doi: 10.1111/geb.12753.

M. Kalamandeen, E. Gloor, E. Mitchard, D. Quincey, G. Ziv, D. Spracklen, B. Spracklen, M. Adami, L. E. O. C. Aragão, and D. Galbraith. Pervasive Rise of Small-scale Deforestation in Amazonia. *Sci Rep*, 8(1):1600, Jan. 2018. ISSN 2045-2322. doi: 10.1038/s41598-018-19358-2.

J. O. Kaplan and H.-K. Lau. The WGLC global gridded monthly lightning stroke density and climatology, July 2019.

E. S. Kasischke and M. R. Turetsky. Recent changes in the fire regime across the North American boreal region—Spatial and temporal patterns of burning across Canada and Alaska. *Geophysical Research Letters*, 33(9), 2006. ISSN 1944-8007. doi: 10.1029/2006GL025677.

R. E. Keane, R. Burgan, and J. van Wagtendonk. Mapping wildland fuels for fire management across multiple scales: Integrating remote sensing, GIS, and biophysical modeling. *Int. J. Wildland Fire*, 10(4):301, 2001. ISSN 1049-8001. doi: 10.1071/WF01028.

J. E. Keeley, C. J. Fotheringham, and M. Morais. Reexamining Fire Suppression Impacts on Brushland Fire Regimes. *Science*, 284(5421):1829–1832, June 1999. doi: 10.1126/science.284.5421.1829.

D. I. Kelley, I. C. Prentice, S. P. Harrison, H. Wang, M. Simard, J. B. Fisher, and K. O. Willis. A comprehensive benchmarking system for evaluating global vegetation models. *Biogeosciences*, 10(5):3313–3340, May 2013. ISSN 1726-4170. doi: 10.5194/bg-10-3313-2013.

D. I. Kelley, I. Bistinas, R. Whitley, C. Burton, T. R. Marthews, and N. Dong. How contemporary bioclimatic and human controls change global fire regimes. *Nat. Clim. Chang.*, 9(9):690–696, Aug. 2019. ISSN 1758-6798. doi: 10.1038/s41558-019-0540-7.

D. I. Kelley, C. Burton, and J. Teixeira. Bringing fire into UKESM, Jan. 2020.

G. Keppel-Aleks, A. S. Wolf, M. Mu, S. C. Doney, D. C. Morton, P. S. Kasibhatla, J. B. Miller, E. J. Dlugokencky, and J. T. Randerson. Separating the influence of temperature, drought, and fire on interannual variability in atmospheric CO2. *Global Biogeochemical Cycles*, 28(11): 1295–1310, 2014. ISSN 1944-9224. doi: 10.1002/2014GB004890.

M. C. Kirchmeier-Young, N. P. Gillett, F. W. Zwiers, A. J. Cannon, and F. S. Anslow. Attribution of the Influence of Human-Induced Climate Change on an Extreme Fire Season. *Earth's Future*, 7(1):2–10, 2019. ISSN 2328-4277. doi: 10.1029/2018EF001050.

C. Klein Goldewijk. Anthropogenic land-use estimates for the Holocene; HYDE 3.2, 2017.

S. Kloster and G. Lasslop. Historical and future fire occurrence (1850 to 2100) simulated in CMIP5 Earth System Models. *Global and Planetary Change*, 150:58–69, Mar. 2017. ISSN 0921-8181. doi: 10.1016/j.gloplacha.2016.12.017.

S. Kloster, N. M. Mahowald, J. T. Randerson, and P. J. Lawrence. The impacts of climate, land use, and demography on fires during the 21st century simulated by CLM-CN. *Biogeosciences*, 9(1):509–525, Jan. 2012. ISSN 1726-4170. doi: 10.5194/bg-9-509-2012.

T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing. Jupyter Notebooks – a publishing format for reproducible computational workflows. In F. Loizides and B. Schmidt, editors, *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, pages 87–90. IOS Press, IOS Press, 2016. doi: 10.3233/978-1-61499-649-1-87.

W. Knorr, T. Kaminski, A. Arneth, and U. Weber. Impact of human population density on fire

frequency at the global scale. *Biogeosciences*, 11(4):1085–1102, Feb. 2014. ISSN 1726-4170. doi: 10.5194/bg-11-1085-2014.

W. Knorr, L. Jiang, and A. Arneth. Climate, CO2 and human population impacts on global wildfire emissions. *Biogeosciences*, 13(1):267–282, Jan. 2016. ISSN 1726-4170. doi: 10.5194/bg-13-267-2016.

P. Köhler, L. Guanter, and J. Joiner. A linear method for the retrieval of sun-induced chlorophyll fluorescence from GOME-2 and SCIAMACHY data. *Atmospheric Measurement Techniques*, 8 (6):2589–2608, June 2015. ISSN 1867-1381. doi: 10.5194/amt-8-2589-2015.

M. A. Krawchuk and M. A. Moritz. Constraints on global fire activity vary across a resource gradient. *Ecology*, 92(1):121–132, Jan. 2011. ISSN 1939-9170. doi: 10.1890/09-1843.1.

M. A. Krawchuk, M. A. Moritz, M.-A. Parisien, J. V. Dorn, and K. Hayhoe. Global Pyrogeography: The Current and Future Distribution of Wildfire. *PLOS ONE*, 4(4):e5102, Apr. 2009. ISSN 1932-6203. doi: 10.1371/journal.pone.0005102.

E. S. Krueger, T. E. Ochsner, J. D. Carlson, D. M. Engle, D. Twidwell, S. D. Fuhlendorf, E. S. Krueger, T. E. Ochsner, J. D. Carlson, D. M. Engle, D. Twidwell, and S. D. Fuhlendorf. Concurrent and antecedent soil moisture relate positively or negatively to probability of large wildfires depending on season. *Int. J. Wildland Fire*, 25(6):657–668, Apr. 2016. ISSN 1448-5516, 1448-5516. doi: 10.1071/WF15104.

A. Kuhn-Régnier. era5analysis (version 0.2.1). Zenodo. https://doi.org/10.5281/zenodo.4173493, Nov. 2020.

A. Kuhn-Régnier. empirical-fire-modelling (version 0.1.2). Zenodo. https://doi.org/10.5281/zenodo.4778777, May 2021.

A. Kuhn-Régnier. jules_output_analysis (version 0.1.0). Zenodo. https://doi.org/10.5281/zenodo.7090997, Sept. 2022a.

A. Kuhn-Régnier. python-inferno (version 1.1.0). Zenodo. https://doi.org/10.5281/zenodo.7113917, Sept. 2022b.

A. Kuhn-Régnier. wildfires (version 1.0.0). Zenodo. https://doi.org/10.5281/zenodo.7091055, Sept. 2022c.

A. Kuhn-Régnier, M. Jumelle, and S. Rajaratnam. ALEPython (version 0.5.5). Zenodo. http://doi.org/10.5281/zenodo.4739201, May 2021a.

A. Kuhn-Régnier, A. Voulgarakis, P. Nowack, M. Forkel, I. C. Prentice, and S. P. Harrison. The importance of antecedent vegetation and drought conditions as global drivers of burnt area. *Biogeosciences*, 18(12):3861–3879, June 2021b. ISSN 1726-4170. doi: 10.5194/bg-18-3861-2021.

A. Kuhn-Régnier, J. Herman, W. Usher, T. Iwanaga, L. Teixeira, N. CELLIER, C. Whealton, D. Hadka, xantares, bernardoct, F. Rios, C. Mutel, E. Cederstrand, T. Andersen, J. van Engelen, ANtlord, H. Kranas, V. K. Dixit, and B. Steubing. SALib v1.5.0c. Zenodo. https://doi.org/10.5281/zenodo.7090976, Sept. 2022a.

A. Kuhn-Régnier, T. Houska, Ben, P. Kraft, H. Agrawal, F. U. Jehn, P. Lauer, Iacopo, A. Vinayak, kbstn, gitporst, J. Castilla-Rho, J-M, Cecile, C. Krause, J. Lee, T. Hallouin, and Y. Wu. SPOTPY v.1.5.16.2. Zenodo. https://doi.org/10.5281/zenodo.7090979, Sept. 2022b.

G. Lasslop, S. Hantson, and S. Kloster. Influence of wind speed on the global variability of burned fraction: A global fire model's perspective. *Int. J. Wildland Fire*, 24(7):989–1000, Aug. 2015. ISSN 1448-5516. doi: 10.1071/WF15052.

G. Lasslop, A. I. Coppola, A. Voulgarakis, C. Yue, and S. Veraverbeke. Influence of Fire on the Carbon Cycle and Climate. *Curr Clim Change Rep*, 5(2):112–123, June 2019. ISSN 2198-6061. doi: 10.1007/s40641-019-00128-9.

G. Lasslop, S. Hantson, S. P. Harrison, D. Bachelet, C. Burton, M. Forkel, M. Forrest, F. Li, J. R. Melton, C. Yue, S. Archibald, S. Scheiter, A. Arneth, T. Hickler, and S. Sitch. Global ecosystems and fire: Multi-model assessment of fire-induced tree cover and carbon storage

reduction. *Glob Change Biol*, page gcb.15160, May 2020. ISSN 1354-1013, 1365-2486. doi: 10.1111/gcb.15160.

Y. Le Page, D. Oom, J. M. N. Silva, P. Jönsson, and J. M. C. Pereira. Seasonality of vegetation fires as modified by human action: Observing the deviation from eco-climatic fire regimes. *Global Ecology and Biogeography*, May 2010. ISSN 1466822X, 14668238. doi: 10.1111/j.1466-8238.2010.00525.x.

Y. Le Page, D. Morton, B. Bond-Lamberty, J. M. C. Pereira, and G. Hurtt. HESFIRE: A global fire model to explore the role of anthropogenic and weather drivers. *Biogeosciences*, 12(3): 887–903, Feb. 2015. ISSN 1726-4189. doi: 10.5194/bg-12-887-2015.

L. A. Lee, K. S. Carslaw, K. J. Pringle, G. W. Mann, and D. V. Spracklen. Emulation of a complex global aerosol model to quantify sensitivity to uncertain parameters. *Atmospheric Chemistry and Physics*, 11(23):12253–12273, Dec. 2011. ISSN 1680-7316. doi: 10.5194/acp-11-12253-2011.

B. A. Leys, J. R. Marlon, C. Umbanhowar, and B. Vannière. Global fire history of grassland biomes. *Ecology and Evolution*, 8(17):8831–8852, 2018. ISSN 2045-7758. doi: 10.1002/ece3.4394.

F. Li, M. Val Martin, M. O. Andreae, A. Arneth, S. Hantson, J. W. Kaiser, G. Lasslop, C. Yue, D. Bachelet, M. Forrest, E. Kluzek, X. Liu, S. Mangeon, J. R. Melton, D. S. Ward, A. Darmenov, T. Hickler, C. Ichoku, B. I. Magi, S. Sitch, G. R. van der Werf, C. Wiedinmyer, and S. S. Rabin. Historical (1700–2012) global multi-model estimates of the fire emissions from the Fire Modeling Intercomparison Project (FireMIP). *Atmos. Chem. Phys.*, 19(19):12545–12567, Oct. 2019. ISSN 1680-7324. doi: 10.5194/acp-19-12545-2019.

J. Li and D. P. Roy. A Global Analysis of Sentinel-2A, Sentinel-2B and Landsat-8 Data Revisit Intervals and Implications for Terrestrial Monitoring. *Remote Sensing*, 9(9):902, Sept. 2017. ISSN 2072-4292. doi: 10.3390/rs9090902.

W. Li, N. MacBean, P. Ciais, P. Defourny, C. Lamarche, S. Bontemps, I. Moreau, R. A. Houghton, and S. Peng. Gross and net land cover changes in the main plant functional types derived

from the annual ESA CCI land cover maps (1992–2015). *Earth System Science Data*, 10(1): 219–234, Jan. 2018. ISSN 1866-3508. doi: 10.5194/essd-10-219-2018.

Z. Li and H. A. Scheraga. Monte Carlo-minimization approach to the multiple-minima problem in protein folding. *Proc Natl Acad Sci U S A*, 84(19):6611–6615, Oct. 1987. ISSN 0027-8424.

J. S. Littell, D. McKenzie, D. L. Peterson, and A. L. Westerling. Climate and wildfire area burned in western U.S. ecoprovinces, 1916–2003. *Ecological Applications*, 19(4):1003–1021, 2009. ISSN 1939-5582. doi: 10.1890/07-1183.1.

S. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. In I. Guyon, R. Fergus, H. Wallach, U. von Luxburg, R. Garnett, S. Vishwanathan, and S. Bengio, editors, *Advances in Neural Information Processing Systems*, volume 2017-December, pages 4766–4775. Neural information processing systems foundation, 2017.

S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1):56–67, Jan. 2020. ISSN 2522-5839. doi: 10.1038/s42256-019-0138-9.

B. I. Magi, S. Rabin, E. Shevliakova, and S. Pacala. Separating agricultural and non-agricultural fire seasonality at regional scales. *Biogeosciences*, 9(8):3003–3012, Aug. 2012. ISSN 1726-4189. doi: 10.5194/bg-9-3003-2012.

S. Mangeon. *Developing and Evaluating a Global Model for Landscape Fires*. Thesis, Imperial College London, London, Feb. 2017.

S. Mangeon, R. Field, M. Fromm, C. McHugh, and A. Voulgarakis. Satellite versus ground-based estimates of burned area: A comparison between MODIS based burned area and fire agency reports over North America in 2007. *The Anthropocene Review*, 3(2):76–92, Aug. 2016a. ISSN 2053-0196. doi: 10.1177/2053019615588790.

S. Mangeon, A. Voulgarakis, R. Gilham, A. Harper, S. Sitch, and G. Folberth. INFERNO: A fire and emissions scheme for the UK Met Office's Unified Model. *Geoscientific Model Development*, 9(8):2685–2700, Aug. 2016b. ISSN 1991-9603. doi: 10.5194/gmd-9-2685-2016.

L. A. Mansfield, P. J. Nowack, M. Kasoar, R. G. Everitt, W. J. Collins, and A. Voulgarakis. Predicting global patterns of long-term climate change from short-term simulations using machine learning. *npj Climate and Atmospheric Science*, 3(1):1–9, Nov. 2020. ISSN 2397-3722. doi: 10.1038/s41612-020-00148-5.

J.-M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder. Approximate Bayesian computational methods. *Stat Comput*, 22(6):1167–1180, Nov. 2012. ISSN 1573-1375. doi: 10.1007/s11222-011-9288-2.

M. E. Marlier, R. S. DeFries, A. Voulgarakis, P. L. Kinney, J. T. Randerson, D. T. Shindell, Y. Chen, and G. Faluvegi. El Niño and health risks from landscape fire emissions in southeast Asia. *Nature Clim Change*, 3(2):131–136, Feb. 2013. ISSN 1758-6798. doi: 10.1038/nclimate1658.

J. R. Marlon, P. J. Bartlein, C. Carcaillet, D. G. Gavin, S. P. Harrison, P. E. Higuera, F. Joos, M. J. Power, and I. C. Prentice. Climate and human influences on global biomass burning over the past two millennia. *Nature Geosci*, 1(10):697–702, Oct. 2008. ISSN 1752-0908. doi: 10.1038/ngeo313.

J. R. Marlon, P. J. Bartlein, D. G. Gavin, C. J. Long, R. S. Anderson, C. E. Briles, K. J. Brown, D. Colombaroli, D. J. Hallett, M. J. Power, E. A. Scharf, and M. K. Walsh. Long-term perspective on wildfires in the western USA. *PNAS*, 109(9):E535–E543, Feb. 2012. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1112839109.

J. R. Marlon, R. Kelly, A.-L. Daniau, B. Vannière, M. J. Power, P. Bartlein, P. Higuera, O. Blarquez, S. Brewer, T. Brücher, A. Feurdean, G. G. Romera, V. Iglesias, S. Y. Maezumi, B. Magi, C. J. Courtney Mustaphi, and T. Zhihai. Reconstructions of biomass burning

from sediment-charcoal records to improve data–model comparisons. *Biogeosciences*, 13(11): 3225–3244, June 2016. ISSN 1726-4170. doi: 10.5194/bg-13-3225-2016.

J. Martínez, C. Vega-Garcia, and E. Chuvieco. Human-caused wildfire risk rating for prevention planning in Spain. *Journal of Environmental Management*, 90(2):1241–1252, Feb. 2009. ISSN 0301-4797. doi: 10.1016/j.jenvman.2008.07.005.

C. Mathison, E. Burke, A. J. Hartley, D. I. Kelley, C. Burton, E. Robertson, N. Gedney, K. Williams, A. Wiltshire, R. J. Ellis, A. Sellar, and C. Jones. Description and Evaluation of the JULES-ES setup for ISIMIP2b. *EGUsphere*, pages 1–24, Dec. 2022. doi: 10.5194/egusphere-2022-1196.

Met Office. *Iris: A Python Library for Analysing and Visualising Meteorological and Oceanographic Data Sets*. Exeter, Devon, v2.4 edition, 2010.

H. Meyer, C. Reudenbach, S. Wöllauer, and T. Nauss. Importance of spatial predictor variable selection in machine learning applications – Moving from data reproduction to spatial prediction. *Ecological Modelling*, 411:108815, Nov. 2019. ISSN 0304-3800. doi: 10.1016/j.ecolmodel.2019.108815.

L. Moesinger, W. Dorigo, R. De Jeu, R. Van der Schalie, T. Scanlon, I. Teubner, and M. Forkel. The Global Long-term Microwave Vegetation Optical Depth Climate Archive VODCA (Version 1.0) [Data set], Feb. 2019.

L. Moesinger, W. Dorigo, R. de Jeu, R. van der Schalie, T. Scanlon, I. Teubner, and M. Forkel. The global long-term microwave Vegetation Optical Depth Climate Archive (VODCA). *Earth System Science Data*, 12(1):177–196, Jan. 2020. ISSN 1866-3508. doi: 10.5194/essd-12-177-2020.

G. H. Mohammed, R. Colombo, E. M. Middleton, U. Rascher, C. van der Tol, L. Nedbal, Y. Goulas, O. Pérez-Priego, A. Damm, M. Meroni, J. Joiner, S. Cogliati, W. Verhoef, Z. Malenovský, J.-P. Gastellu-Etchegorry, J. R. Miller, L. Guanter, J. Moreno, I. Moya, J. A. Berry, C. Frankenberg, and P. J. Zarco-Tejada. Remote sensing of solar-induced chlorophyll fluorescence (SIF) in

vegetation: 50 years of progress. *Remote Sensing of Environment*, 231:111177, Sept. 2019. ISSN 0034-4257. doi: 10.1016/j.rse.2019.04.030.

C. Molnar. *Interpretable Machine Learning*. Lulu Press, Morrisville, North Carolina, United States, Apr. 2020. ISBN 978-0-244-76852-2.

B. Moreira, J. Tormo, E. Estrelles, and J. G. Pausas. Disentangling the role of heat and smoke as germination cues in Mediterranean Basin flora. *Ann Bot*, 105(4):627–635, Apr. 2010. ISSN 0305-7364. doi: 10.1093/aob/mcq017.

P. Morgan, C. C. Hardy, T. W. Swetnam, M. G. Rollins, and D. G. Long. Mapping fire regimes across time and space: Understanding coarse and fine-scale fire patterns. *Int. J. Wildland Fire*, 10(4):329, 2001. ISSN 1049-8001. doi: 10.1071/WF01032.

M. A. Moritz. Spatiotemporal Analysis of Controls on Shrubland Fire Regimes: Age Dependency and Fire Hazard. *Ecology*, 84(2):351–361, 2003. ISSN 1939-9170. doi: 10.1890/0012-9658(2003)084[0351:SAOCOS]2.0.CO;2.

M. A. Moritz, M.-A. Parisien, E. Batllori, M. A. Krawchuk, J. Van Dorn, D. J. Ganz, and K. Hayhoe. Climate change and disruptions to global fire activity. *Ecosphere*, 3(6):art49, June 2012. ISSN 2150-8925. doi: 10.1890/ES11-00345.1.

D. C. Morton, R. S. DeFries, Y. E. Shimabukuro, L. O. Anderson, E. Arai, F. del Bon Espirito-Santo, R. Freitas, and J. Morisette. Cropland expansion changes deforestation dynamics in the southern Brazilian Amazon. *Proc. Natl. Acad. Sci. U.S.A.*, 103(39):14637–14641, Sept. 2006. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0606377103.

R. Myneni, Y. Knyazikhin, and T. Park. MOD15A2H MODIS/Terra Leaf Area Index/FPAR 8-Day L4 Global 500m SIN Grid V006, 2015.

R. B. Myneni and D. L. Williams. On the relationship between FAPAR and NDVI. *Remote Sensing of Environment*, 49(3):200–211, Sept. 1994. ISSN 0034-4257. doi: 10.1016/0034-4257(94)90016-7.

J. E. Nash and J. V. Sutcliffe. River flow forecasting through conceptual models part I — A discussion of principles. *Journal of Hydrology*, 10(3):282–290, Apr. 1970. ISSN 0022-1694. doi: 10.1016/0022-1694(70)90255-6.

J. A. Nelder and R. W. M. Wedderburn. Generalized Linear Models. *Journal of the Royal Statistical Society. Series A (General)*, 135(3):370–384, 1972. ISSN 0035-9238. doi: 10.2307/2344614.

P. Noojipady, C. D. Morton, N. M. Macedo, C. D. Victoria, C. Huang, K. H. Gibbs, and L. E. Bolfe. Forest carbon emissions from cropland expansion in the Brazilian Cerrado biome. *Environ. Res. Lett.*, 12(2):025004, Feb. 2017. ISSN 1748-9326. doi: 10.1088/1748-9326/aa5986.

J. Norton. An introduction to sensitivity assessment of simulation models. *Environmental Modelling & Software*, 69:166–174, July 2015. ISSN 1364-8152. doi: 10.1016/j.envsoft.2015.03.020.

P. Nowack, P. Braesicke, J. Haigh, N. L. Abraham, J. Pyle, and A. Voulgarakis. Using machine learning to build temperature-based ozone parameterizations for climate sensitivity simulations. *Environ. Res. Lett.*, 13(10):104016, Oct. 2018. ISSN 1748-9326. doi: 10.1088/1748-9326/aae2be.

P. Nowack, J. Runge, V. Eyring, and J. D. Haigh. Causal networks for climate model evaluation and constrained projections. *Nature Communications*, 11(1):1415, Mar. 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-15195-y.

S. O, X. Hou, and R. Orth. Observational evidence of wildfire-promoting soil moisture anomalies. *Scientific Reports*, 10(1):11008, July 2020. ISSN 2045-2322. doi: 10.1038/s41598-020-67530-4.

B. O. Ogutu, J. Dash, and T. P. Dawson. Evaluation of the influence of two operational fraction of absorbed photosynthetically active radiation (FAPAR) products on terrestrial ecosystem productivity modelling. *International Journal of Remote Sensing*, 35(1):321–340, Jan. 2014. ISSN 0143-1161. doi: 10.1080/01431161.2013.871083.

T. E. Oliphant. *A Guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.

B. Olson, I. Hashmi, K. Molloy, and A. Shehu. Basin Hopping as a General and Versatile Optimization Framework for the Characterization of Biological Macromolecules. *Advances in Artificial Intelligence*, 2012:e674832, Dec. 2012. ISSN 1687-7470. doi: 10.1155/2012/674832.

F. Oris, H. Asselin, A. A. Ali, W. Finsinger, and Y. Bergeron. Effect of increased fire activity on global warming in the boreal forest. *Environ. Rev.*, 22(3):206–219, Sept. 2014. ISSN 1181-8700. doi: 10.1139/er-2013-0062.

M. Owe, R. de Jeu, and J. Walker. A methodology for surface soil moisture and vegetation optical depth retrieval using the microwave polarization difference index. *IEEE Transactions on Geoscience and Remote Sensing*, 39(8):1643–1654, Aug. 2001. ISSN 1558-0644. doi: 10.1109/36.942542.

A. B. Owen. Sobol' Indices and Shapley Value. *SIAM/ASA J. Uncertainty Quantification*, 2(1): 245–251, Jan. 2014. doi: 10.1137/130936233.

M. Padilla, S. Stehman, J. Litago, and E. Chuvieco. Assessing the Temporal Stability of the Accuracy of a Time Series of Burned Area Products. *Remote Sensing*, 6(3):2050–2068, Mar. 2014a. ISSN 2072-4292. doi: 10.3390/rs6032050.

M. Padilla, S. V. Stehman, and E. Chuvieco. Validation of the 2008 MODIS-MCD45 global burned area product using stratified random sampling. *Remote Sensing of Environment*, 144: 187–196, Mar. 2014b. ISSN 0034-4257.

M. Padilla, S. V. Stehman, R. Ramo, D. Corti, S. Hantson, P. Oliva, I. Alonso-Canas, A. V. Bradley, K. Tansey, B. Mota, J. M. Pereira, and E. Chuvieco. Comparing the accuracies of remote sensing global burned area products using stratified random sampling and estimation. *Remote Sensing of Environment*, 160:114–121, Apr. 2015. ISSN 0034-4257. doi: 10.1016/j.rse.2015.01.005.

S. A. Parks, C. Miller, M.-A. Parisien, L. M. Holsinger, S. Z. Dobrowski, and J. Abatzoglou.

Wildland fire deficit and surplus in the western United States, 1984–2012. *Ecosphere*, 6(12): 1–13, Dec. 2015. ISSN 2150-8925. doi: 10.1890/ES15-00294.1.

S. Paula and J. G. Pausas. Burning seeds: Germinative response to heat treatments in relation to resprouting ability: Post-fire germination and resprouting. *Journal of Ecology*, 96(3):543–552, Feb. 2008. ISSN 00220477. doi: 10.1111/j.1365-2745.2008.01359.x.

J. G. Pausas and W. J. Bond. Alternative Biome States in Terrestrial Ecosystems. *Trends in Plant Science*, 25(3):250–263, Mar. 2020. ISSN 1360-1385. doi: 10.1016/j.tplants.2019.11.003.

J. G. Pausas, B. B. Lamont, S. Paula, B. Appezzato-da-Glória, and A. Fidelis. Unearthing belowground bud banks in fire-prone ecosystems. *New Phytologist*, 217(4):1435–1448, 2018. ISSN 1469-8137. doi: 10.1111/nph.14982.

O. Pechony and D. T. Shindell. Fire parameterization on a global scale. *Journal of Geophysical Research: Atmospheres*, 114(D16), 2009. ISSN 2156-2202. doi: 10.1029/2009JD011927.

O. Pechony and D. T. Shindell. Driving forces of global wildfires over the past millennium and the forthcoming century. *Proceedings of the National Academy of Sciences*, 107(45):19167–19170, Nov. 2010. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1003669107.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

O. Perkins, C. Smith, and J. Millington. Human-Fire Interactions: A Global Database, Apr. 2021.

G. L. Perry. Current approaches to modelling the spread of wildland fire: A review. *Progress in Physical Geography: Earth and Environment*, 22(2):222–245, June 1998. ISSN 0309-1333, 1477-0296. doi: 10.1177/030913339802200204.

M. L. Pettinari and E. Chuvieco. Generation of a global fuel data set using the Fuel Characteristic Classification System. *Biogeosciences*, 13(7):2061–2076, Apr. 2016. ISSN 1726-4189. doi: 10.5194/bg-13-2061-2016.

M. Pfeiffer, A. Spessa, and J. O. Kaplan. A model for global biomass burning in preindustrial time: LPJ-LMfire (v1.0). *Geosci. Model Dev.*, 6(3):643–685, May 2013. ISSN 1991-9603. doi: 10.5194/gmd-6-643-2013.

P. Ploton, F. Mortier, M. Réjou-Méchain, N. Barbier, N. Picard, V. Rossi, C. Dormann, G. Cornu, G. Viennois, N. Bayol, A. Lyapustin, S. Gourlet-Fleury, and R. Pélissier. Spatial validation reveals poor predictive performance of large-scale ecological mapping models. *Nature Communications*, 11(1):4540, Sept. 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-18321-y.

S. Plummer, O. Arino, M. Simon, and W. Steffen. Establishing A Earth Observation Product Service For The Terrestrial Carbon Community: The Globcarbon Initiative. *Mitig Adapt Strat Glob Change*, 11(1):97–111, Jan. 2006. ISSN 1573-1596. doi: 10.1007/s11027-006-1012-8.

B. Poulter, N. MacBean, A. Hartley, I. Khlystova, O. Arino, R. Betts, S. Bontemps, M. Boettcher, C. Brockmann, P. Defourny, S. Hagemann, M. Herold, G. Kirches, C. Lamarche, D. Lederer, C. Ottlé, M. Peters, and P. Peylin. Plant functional type classification for earth system models: Results from the European Space Agency's Land Cover Climate Change Initiative. *Geoscientific Model Development*, 8(7):2315–2328, July 2015. ISSN 1991-959X. doi: 10.5194/gmd-8-2315-2015.

I. C. Prentice. The Burning Issue. *Science*, 330(6011):1636–1637, Dec. 2010. doi: 10.1126/science.1199809.

P. Probst, M. N. Wright, and A.-L. Boulesteix. Hyperparameters and tuning strategies for random forest. *WIREs Data Mining and Knowledge Discovery*, 9(3):e1301, 2019. ISSN 1942-4795. doi: 10.1002/widm.1301.

S. S. Rabin, J. R. Melton, G. Lasslop, D. Bachelet, M. Forrest, S. Hantson, J. O. Kaplan, F. Li,

S. Mangeon, D. S. Ward, C. Yue, V. K. Arora, T. Hickler, S. Kloster, W. Knorr, L. Nieradzik, A. Spessa, G. A. Folberth, T. Sheehan, A. Voulgarakis, D. I. Kelley, I. C. Prentice, S. Sitch, S. Harrison, and A. Arneth. The Fire Modeling Intercomparison Project (FireMIP), phase 1: Experimental and analytical protocols with detailed model descriptions. *Geoscientific Model Development*, 10(3):1175–1197, Mar. 2017. ISSN 1991-959X. doi: 10.5194/gmd-10-1175-2017.

J. T. Randerson, G. R. van der Werf, G. J. Collatz, L. Giglio, C. J. Still, P. Kasibhatla, J. B. Miller, J. W. C. White, R. S. DeFries, and E. S. Kasischke. Fire emissions from C3 and C4 vegetation and their influence on interannual variability of atmospheric CO2 and δ13CO2. *Global Biogeochemical Cycles*, 19(2), May 2005. ISSN 1944-9224. doi: 10.1029/2004GB002366.

J. T. Randerson, Y. Chen, G. R. van der Werf, B. M. Rogers, and D. C. Morton. Global burned area and biomass burning emissions from small fires. *Journal of Geophysical Research: Biogeosciences*, 117(G4), 2012. ISSN 2156-2202. doi: 10.1029/2012JG002128.

J. Ratnam, W. J. Bond, R. J. Fensham, W. A. Hoffmann, S. Archibald, C. E. R. Lehmann, M. T. Anderson, S. I. Higgins, and M. Sankaran. When is a 'forest' a savanna, and why does it matter? *Global Ecology and Biogeography*, 20(5):653–660, 2011. ISSN 1466-8238. doi: 10.1111/j.1466-8238.2010.00634.x.

B. Ripley, V. Visser, P.-A. Christin, S. Archibald, T. Martin, and C. Osborne. Fire ecology of C3 and C4 grasses depends on evolutionary history and frequency of burning but not photosynthetic type. *Ecology*, 96(10):2679–2691, 2015. ISSN 1939-9170. doi: 10.1890/14-1495.1.

D. R. Roberts, V. Bahn, S. Ciuti, M. S. Boyce, J. Elith, G. Guillera-Arroita, S. Hauenstein, J. J. Lahoz-Monfort, B. Schröder, W. Thuiller, D. I. Warton, B. A. Wintle, F. Hartig, and C. F. Dormann. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8):913–929, Aug. 2017. ISSN 1600-0587. doi: 10.1111/ecog.02881.

G. Roberts and M. J. Wooster. Global impact of landscape fire emissions on surface level PM2.5

concentrations, air quality exposure and population mortality. *Atmospheric Environment*, 252: 118210, May 2021. ISSN 1352-2310. doi: 10.1016/j.atmosenv.2021.118210.

G. Roberts, M. J. Wooster, N. Lauret, J. P. Gastellu-Etchegorry, T. Lynham, and D. McRae. Investigating the impact of overlying vegetation canopy structures on fire radiative power (FRP) retrieval through simulation and measurement. *Remote Sensing of Environment*, 217: 158–171, Nov. 2018. ISSN 0034-4257. doi: 10.1016/j.rse.2018.08.015.

E. Robinson, G. E. G. Russell, W. Trollope, and B. Downing. Short-term burning treatments and ecological interactions in the herb layer of false thornveld of the eastern province. *Proceedings of the Annual Congresses of the Grassland Society of Southern Africa*, 14(1):79–83, Jan. 1979. ISSN 0072-5560. doi: 10.1080/00725560.1979.9648864.

C. J. Rodger, J. B. Brundell, R. L. Dowden, and N. R. Thomson. Location accuracy of long distance VLF lightning locationnetwork. *Annales Geophysicae*, 22(3):747–758, Mar. 2004. ISSN 0992-7689. doi: 10.5194/angeo-22-747-2004.

E. Roteta, A. Bastarrika, M. Padilla, T. Storm, and E. Chuvieco. Development of a Sentinel-2 burned area algorithm: Generation of a small fire database for sub-Saharan Africa. *Remote Sensing of Environment*, 222:1–17, Mar. 2019. ISSN 00344257. doi: 10.1016/j.rse.2018.12.011.

D. P. Roy and L. Boschetti. Southern Africa Validation of the MODIS, L3JRC, and GlobCarbon Burned-Area Products. *IEEE Transactions on Geoscience and Remote Sensing*, 47(4):1032–1044, Apr. 2009. ISSN 1558-0644. doi: 10.1109/TGRS.2008.2009000.

D. P. Roy, L. Boschetti, C. O. Justice, and J. Ju. The collection 5 MODIS burned area product — Global evaluation by comparison with the MODIS active fire product. *Remote Sensing of Environment*, 112(9):3690–3707, Sept. 2008. ISSN 0034-4257. doi: 10.1016/j.rse.2008.05.013.

D. P. Roy, H. Huang, L. Boschetti, L. Giglio, L. Yan, H. H. Zhang, and Z. Li. Landsat-8 and Sentinel-2 burned area mapping - A combined sensor multi-temporal change detection

approach. *Remote Sensing of Environment*, 231:111254, Sept. 2019. ISSN 0034-4257. doi: 10.1016/j.rse.2019.111254.

J. Runge, P. Nowack, M. Kretschmer, S. Flaxman, and D. Sejdinovic. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5(11):eaau4996, Nov. 2019. ISSN 2375-2548. doi: 10.1126/sciadv.aau4996.

Y. Ryu, J. A. Berry, and D. D. Baldocchi. What is global photosynthesis? History, uncertainties and opportunities. *Remote Sensing of Environment*, 223:95–114, Mar. 2019. ISSN 0034-4257. doi: 10.1016/j.rse.2019.01.016.

M. Sadegh and J. A. Vrugt. Bridging the gap between GLUE and formal statistical approaches: Approximate Bayesian computation. *Hydrology and Earth System Sciences*, 17(12):4831–4850, Dec. 2013. ISSN 1027-5606. doi: 10.5194/hess-17-4831-2013.

M. Sadegh and J. A. Vrugt. Approximate Bayesian Computation using Markov Chain Monte Carlo simulation: DREAM(ABC). *Water Resources Research*, 50(8):6767–6787, Aug. 2014. ISSN 1944-7973. doi: 10.1002/2014WR015386.

A. Saltelli. Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145(2):280–297, May 2002. ISSN 0010-4655. doi: 10.1016/S0010-4655(02)00280-1.

A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global Sensitivity Analysis. The Primer*. John Wiley & Sons, Ltd, Chichester, UK, 2008. ISBN 978-0-470-72518-4 978-0-470-05997-5. doi: 10.1002/9780470725184.

A. Saltelli, P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Computer Physics Communications*, 181(2):259–270, Feb. 2010. ISSN 0010-4655. doi: 10.1016/j.cpc.2009.09.018.

B. M. Sanderson and R. A. Fisher. A fiery wake-up call for climate science. *Nat. Clim. Chang.*, 10(3):175–177, Mar. 2020. ISSN 1758-678X, 1758-6798. doi: 10.1038/s41558-020-0707-2.

R. C. Scholten, R. Jandt, E. A. Miller, B. M. Rogers, and S. Veraverbeke. Overwintering fires in boreal forests. *Nature*, 593(7859):399–404, May 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03437-y.

N. H. Schumaker, S. M. Watkins, and J. A. Heinrichs. HexFire: A Flexible and Accessible Wildfire Simulator. *Land*, 11(8):1288, Aug. 2022. ISSN 2073-445X. doi: 10.3390/land11081288.

A. C. Scott. *Burning Planet: The Story of Fire through Time*. Oxford University Press, Oxford, UK, first edition edition, 2018. ISBN 978-0-19-873484-0.

M. R. Segal. Machine Learning Benchmarks and Random Forest Regression. Apr. 2004.

J. Settele and R. Scholes. Terrestrial and inland water systems. In *Climate Change 2014: Impacts,Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. 2014.

M. Simon, S. Plummer, F. Fierens, J. J. Hoelzemann, and O. Arino. Burnt area detection at global scale using ATSR-2: The GLOBSCAR products and their qualification. *Journal of Geophysical Research: Atmospheres*, 109(D14), 2004. ISSN 2156-2202. doi: 10.1029/2003JD003622.

K. J. Simpson, J. K. Olofsson, B. S. Ripley, and C. P. Osborne. Frequent fires prime plant developmental responses to burning. *Proceedings of the Royal Society B: Biological Sciences*, 286(1909):20191315, Aug. 2019. doi: 10.1098/rspb.2019.1315.

S. Sitch, B. Smith, I. C. Prentice, A. Arneth, A. Bondeau, W. Cramer, J. O. Kaplan, S. Levis, W. Lucht, M. T. Sykes, K. Thonicke, and S. Venevsky. Evaluation of ecosystem dynamics, plant geography and terrestrial carbon cycling in the LPJ dynamic global vegetation model: LPJ DYNAMIC GLOBAL VEGETATION MODEL. *Global Change Biology*, 9(2):161–185, Feb. 2003. ISSN 13541013. doi: 10.1046/j.1365-2486.2003.00569.x.

I. M. Sobol′. Sensitivity Estimates for Nonlinear Mathematical Models. *Math. Modell. Comput. Exp.*, 1:407–414, 1993.

I. M. Sobol′. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1):271–280, Feb. 2001. ISSN 0378-4754. doi: 10.1016/S0378-4754(00)00270-6.

W. T. Sommers, R. A. Loehman, and C. C. Hardy. Wildland fire emissions, carbon, and climate: Science overview and knowledge needs. *Forest Ecology and Management*, 317:1–8, Apr. 2014. ISSN 0378-1127. doi: 10.1016/j.foreco.2013.12.014.

A. Spessa, B. McBeth, and C. Prentice. Relationships among fire frequency, rainfall and vegetation patterns in the wet–dry tropics of northern Australia: An analysis based on NOAA-AVHRR data. *Global Ecology and Biogeography*, 14(5):439–454, Aug. 2005. ISSN 1466-8238. doi: 10.1111/j.1466-822x.2005.00174.x.

A. C. Staver, S. Archibald, and S. A. Levin. The Global Extent and Determinants of Savanna and Forest as Alternative Biome States. *Science*, 334(6053):230–232, Oct. 2011. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1210465.

T. W. Swetnam and J. L. Betancourt. Mesoscale Disturbance and Ecological Response to Decadal Climatic Variability in the American Southwest. *Journal of Climate*, 11(12):3128–3147, Dec. 1998. ISSN 0894-8755, 1520-0442. doi: 10.1175/1520-0442(1998)011⟨3128:MDAERT⟩2.0.CO;2.

A. D. Syphard, V. C. Radeloff, J. E. Keeley, T. J. Hawbaker, M. K. Clayton, S. I. Stewart, and R. B. Hammer. Human Influence on California Fire Regimes. *Ecological Applications*, 17(5): 1388–1402, 2007. ISSN 1939-5582. doi: 10.1890/06-1128.1.

K. Tansey, J.-M. Grégoire, D. Stroppiana, A. Sousa, J. Silva, J. M. C. Pereira, L. Boschetti, M. Maggi, P. A. Brivio, R. Fraser, S. Flasse, D. Ershov, E. Binaghi, D. Graetz, and P. Peduzzi. Vegetation burning in the year 2000: Global burned area estimates from SPOT VEGETATION

data. *Journal of Geophysical Research: Atmospheres*, 109(D14), 2004. ISSN 2156-2202. doi: 10.1029/2003JD003598.

K. Tansey, J.-M. Grégoire, P. Defourny, R. Leigh, J.-F. Pekel, E. van Bogaert, and E. Bartholomé. A new, global, multi-annual (2000–2007) burnt area product at 1 km resolution. *Geophysical Research Letters*, 35(1), 2008. ISSN 1944-8007. doi: 10.1029/2007GL031567.

L. Teckentrup, S. P. Harrison, S. Hantson, A. Heil, J. R. Melton, M. Forrest, F. Li, C. Yue, A. Arneth, T. Hickler, S. Sitch, and G. Lasslop. Response of simulated burned area to historical changes in environmental and anthropogenic factors: A comparison of seven fire models. *Biogeosciences*, 16(19):3883–3910, Oct. 2019. ISSN 1726-4189. doi: 10.5194/bg-16-3883-2019.

J. C. Teixeira, G. A. Folberth, F. M. O'Connor, N. Unger, and A. Voulgarakis. Coupling interactive fire with atmospheric composition and climate in the UK Earth System Model. *Geoscientific Model Development*, 14(10):6515–6539, Oct. 2021. ISSN 1991-959X. doi: 10.5194/gmd-14-6515-2021.

I. E. Teubner, M. Forkel, M. Jung, Y. Y. Liu, D. G. Miralles, R. Parinussa, R. van der Schalie, M. Vreugdenhil, C. R. Schwalm, G. Tramontana, G. Camps-Valls, and W. A. Dorigo. Assessing the relationship between microwave vegetation optical depth and gross primary production. *International Journal of Applied Earth Observation and Geoinformation*, 65:79–91, Mar. 2018. ISSN 03032434. doi: 10.1016/j.jag.2017.10.006.

P. B. Thomas, P. J. Watson, R. A. Bradstock, T. D. Penman, and O. F. Price. Modelling surface fine fuel dynamics across climate gradients in eucalypt forests of south-eastern Australia. *Ecography*, 37(9):827–837, 2014. ISSN 1600-0587. doi: 10.1111/ecog.00445.

K. Thonicke, S. Venevsky, S. Sitch, and W. Cramer. The Role of Fire Disturbance for Global Vegetation Dynamics: Coupling Fire into a Dynamic Global Vegetation Model. *Global Ecology and Biogeography*, 10(6):661–677, 2001. ISSN 1466-822X.

M. Thurner, C. Beer, M. Santoro, N. Carvalhais, T. Wutzler, D. Schepaschenko, A. Shvidenko,

E. Kompter, B. Ahrens, S. R. Levick, and C. Schmullius. Carbon stock and density of northern boreal and temperate forests. *Global Ecology and Biogeography*, 23(3):297–310, 2014. ISSN 1466-8238. doi: 10.1111/geb.12125.

D. Tix and I. Charvat. Aboveground Biomass Removal by Burning and Raking Increases Diversity in a Reconstructed Prairie. *Restoration Ecology*, 13(1):20–28, 2005. ISSN 1526-100X. doi: 10.1111/j.1526-100X.2005.00004.x.

M. Turco, J. J. Rosa-Cánovas, J. Bedia, S. Jerez, J. P. Montávez, M. C. Llasat, and A. Provenzale. Exacerbated fires in Mediterranean Europe due to anthropogenic warming projected with non-stationary climate-fire models. *Nat Commun*, 9(1):3821, Dec. 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-06358-z.

M. R. Turetsky, E. S. Kane, J. W. Harden, R. D. Ottmar, K. L. Manies, E. Hoy, and E. S. Kasischke. Recent acceleration of biomass burning and carbon losses in Alaskan forests and peatlands. *Nature Geosci*, 4(1):27–31, Jan. 2011. ISSN 1752-0894, 1752-0908. doi: 10.1038/ngeo1027.

W. Usher, J. Herman, T. Iwanaga, L. Teixeira, N. CELLIER, C. Whealton, D. Hadka, xantares, bernardoct, F. Rios, C. Mutel, E. Cederstrand, TobiasKAndersen, J. van Engelen, ANtlord, H. Kranas, J. Foramitti, S. Menegon, V. K. Dixit, and bsteubing. SALib/SALib: V1.3.13. Zenodo. https://doi.org/10.5281/zenodo.4707524, Apr. 2021.

G. R. van der Werf, J. T. Randerson, L. Giglio, N. Gobron, and A. J. Dolman. Climate controls on the variability of fires in the tropics and subtropics. *Global Biogeochemical Cycles*, 22(3), 2008. ISSN 1944-9224. doi: 10.1029/2007GB003122.

G. R. van der Werf, J. T. Randerson, L. Giglio, T. T. van Leeuwen, Y. Chen, B. M. Rogers, M. Mu, M. J. E. van Marle, D. C. Morton, G. J. Collatz, R. J. Yokelson, and P. S. Kasibhatla. Global fire emissions estimates during 1997–2016. *Earth System Science Data*, 9(2):697–720, Sept. 2017. ISSN 1866-3516. doi: 10.5194/essd-9-697-2017.

P. van Lierop, E. Lindquist, S. Sathyapala, and G. Franceschini. Global forest area disturbance from fire, insect pests, diseases and severe weather events. *Forest Ecology and Management*, 352:78–88, Sept. 2015. ISSN 0378-1127. doi: 10.1016/j.foreco.2015.06.010.

M. J. E. van Marle, G. R. van der Werf, R. a. M. de Jeu, and Y. Y. Liu. Annual South American forest loss estimates based on passive microwave remote sensing (1990–2010). *Biogeosciences*, 13(2):609–624, Feb. 2016. ISSN 1726-4170. doi: 10.5194/bg-13-609-2016.

M. J. E. van Marle, R. D. Field, G. R. van der Werf, I. A. Estrada de Wagt, R. A. Houghton, L. V. Rizzo, P. Artaxo, and K. Tsigaridis. Fire and deforestation dynamics in Amazonia (1973–2014). *Global Biogeochemical Cycles*, 31(1):24–38, Jan. 2017a. ISSN 1944-9224. doi: 10.1002/2016GB005445.

M. J. E. van Marle, S. Kloster, B. I. Magi, J. R. Marlon, A.-L. Daniau, R. D. Field, A. Arneth, M. Forrest, S. Hantson, N. M. Kehrwald, W. Knorr, G. Lasslop, F. Li, S. Mangeon, C. Yue, J. W. Kaiser, and G. R. van der Werf. Historic global biomass burning emissions for CMIP6 (BB4CMIP) based on merging satellite observations with proxies and fire models (1750–2015). *Geoscientific Model Development*, 10(9):3329–3357, Sept. 2017b. ISSN 1991-959X. doi: 10.5194/gmd-10-3329-2017.

E. H. van Nes, A. Staal, S. Hantson, M. Holmgren, S. Pueyo, R. E. Bernardi, B. M. Flores, C. Xu, and M. Scheffer. Fire forbids fifty-fifty forest. *PLOS ONE*, 13(1):e0191027, Jan. 2018. ISSN 1932-6203. doi: 10.1371/journal.pone.0191027.

G. J. van Oldenborgh, F. Krikken, S. Lewis, N. J. Leach, F. Lehner, K. R. Saunders, M. van Weele, K. Haustein, S. Li, D. Wallom, S. Sparrow, J. Arrighi, R. P. Singh, M. K. van Aalst, S. Y. Philip, R. Vautard, and F. E. L. Otto. Attribution of the Australian bushfire risk to anthropogenic climate change. *Natural Hazards and Earth System Sciences Discussions*, pages 1–46, Mar. 2020. ISSN 1561-8633. doi: 10.5194/nhess-2020-69.

G. Van Rossum and F. L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009. ISBN 1-4414-1269-7.

B. W. Van Wilgen, H. Biggs, S. P. O'Regan, and N. Mare. Fire history of the savanna ecosystems in the Kruger National Park, South Africa, between 1941 and 1996. *South African Journal of Science*, 96(4):167–178, Apr. 2000. ISSN 0038-2353.

E. M. Veenendaal, M. Torello-Raventos, H. S. Miranda, N. M. Sato, I. Oliveras, F. van Langevelde, G. P. Asner, and J. Lloyd. On the relationship between fire regime and vegetation structure in the tropics. *New Phytologist*, 218(1):153–166, 2018. ISSN 1469-8137. doi: 10.1111/nph.14940.

S. Venevsky, K. Thonicke, S. Sitch, and W. Cramer. Simulating fire regimes in human-dominated ecosystems: Iberian Peninsula case study. *Global Change Biology*, 8(10):984–998, 2002. ISSN 1365-2486. doi: 10.1046/j.1365-2486.2002.00528.x.

P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, İ. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. . . Contributors. SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, 17(3):261–272, 2020. doi: 10.1038/s41592-019-0686-2.

A. Voulgarakis and R. D. Field. Fire Influences on Atmospheric Composition, Air Quality and Climate. *Curr Pollution Rep*, 1(2):70–81, May 2015. ISSN 2198-6592. doi: 10.1007/s40726-015-0007-z.

A. Voulgarakis, M. E. Marlier, G. Faluvegi, D. T. Shindell, K. Tsigaridis, and S. Mangeon. Interannual variability of tropospheric trace gases and aerosols: The role of biomass burning emissions. *Journal of Geophysical Research: Atmospheres*, 120(14):7157–7173, June 2015. ISSN 2169-8996. doi: 10.1002/2014JD022926.

J. A. Vrugt. Markov chain Monte Carlo simulation using the DREAM software package: Theory, concepts, and MATLAB implementation. *Environmental Modelling & Software*, 75:273–316, Jan. 2016. ISSN 1364-8152. doi: 10.1016/j.envsoft.2015.08.013.

J. A. Vrugt, C. J. F. ter Braak, M. P. Clark, J. M. Hyman, and B. A. Robinson. Treatment of input uncertainty in hydrologic modeling: Doing hydrology backward with Markov chain Monte Carlo simulation. *Water Resources Research*, 44(12), Dec. 2008. ISSN 1944-7973. doi: 10.1029/2007WR006720.

J. A. Vrugt, C. ter Braak, C. Diks, B. A. Robinson, J. M. Hyman, and D. Higdon. Accelerating Markov Chain Monte Carlo Simulation by Differential Evolution with Self-Adaptive Randomized Subspace Sampling. *International Journal of Nonlinear Sciences and Numerical Simulation*, 10(3):273–290, Mar. 2009. ISSN 2191-0294. doi: 10.1515/IJNSNS.2009.10.3.273.

W. Wagner, G. Lemoine, and H. Rott. A Method for Estimating Soil Moisture from ERS Scatterometer and Soil Data. *Remote Sensing of Environment*, 70(2):191–207, Nov. 1999. ISSN 0034-4257. doi: 10.1016/S0034-4257(99)00036-X.

D. J. Wales. *Energy Landscapes*. Cambridge Molecular Science. Cambridge University Press, Cambridge, UK ; New York, 2003. ISBN 978-0-521-81415-7.

D. J. Wales and J. P. K. Doye. Global Optimization by Basin-Hopping and the Lowest Energy Structures of Lennard-Jones Clusters Containing up to 110 Atoms. *J. Phys. Chem. A*, 101 (28):5111–5116, July 1997. ISSN 1089-5639, 1520-5215. doi: 10.1021/jp970984n.

D. J. Wales and H. A. Scheraga. Global Optimization of Clusters, Crystals, and Biomolecules. *Science*, 285(5432):1368–1372, Aug. 1999. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.285.5432.1368.

Z. Wang, J. Chappellaz, K. Park, and J. E. Mak. Large Variations in Southern Hemisphere Biomass Burning During the Last 650 Years. *Science*, 330(6011):1663–1666, Dec. 2010. doi: 10.1126/science.1197257.

D. S. Ward, S. Kloster, N. M. Mahowald, B. M. Rogers, J. T. Randerson, and P. G. Hess. The changing radiative forcing of fires: Global model estimates for past, present and future.

Atmospheric Chemistry and Physics, 12(22):10857–10886, Nov. 2012. ISSN 1680-7316. doi: 10.5194/acp-12-10857-2012.

A. L. Westerling. Warming and Earlier Spring Increase Western U.S. Forest Wildfire Activity. *Science*, 313(5789):940–943, Aug. 2006. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1128834.

A. L. Westerling, A. Gershunov, T. J. Brown, D. R. Cayan, and M. D. Dettinger. Climate and Wildfire in the Western United States. *Bulletin of the American Meteorological Society*, 84(5): 595–604, May 2003. ISSN 0003-0007, 1520-0477. doi: 10.1175/BAMS-84-5-595.

C. Whitlock, P. E. Higuera, D. B. McWethy, and C. E. Briles. Paleoecological Perspectives on Fire Ecology: Revisiting the Fire-Regime Concept. *The Open Ecology Journal*, 3(1), Mar. 2010.

P. Wiederkehr. *Global Sensitivity Analysis with Dependent Inputs*. MSc, ETH Zurich, Zurich, Switzerland, 2018.

O. Wild, A. Voulgarakis, F. O'Connor, J.-F. Lamarque, E. M. Ryan, and L. Lee. Global sensitivity analysis of chemistry–climate model budgets of tropospheric ozone and OH: Exploring model diversity. *Atmospheric Chemistry and Physics*, 20(7):4047–4058, Apr. 2020. ISSN 1680-7316. doi: 10.5194/acp-20-4047-2020.

M. J. Wooster, G. J. Roberts, L. Giglio, D. P. Roy, P. H. Freeborn, L. Boschetti, C. Justice, C. Ichoku, W. Schroeder, D. Davies, A. M. S. Smith, A. Setzer, I. Csiszar, T. Strydom, P. Frost, T. Zhang, W. Xu, M. C. de Jong, J. M. Johnston, L. Ellison, K. Vadrevu, A. M. Sparks, H. Nguyen, J. McCarty, V. Tanpipat, C. Schmidt, and J. San-Miguel-Ayanz. Satellite remote sensing of active fires: History and current status, applications and future requirements. *Remote Sensing of Environment*, 267:112694, Dec. 2021. ISSN 0034-4257. doi: 10.1016/j.rse.2021.112694.

C. Wu, S. Sitch, C. Huntingford, L. M. Mercado, S. Venevsky, G. Lasslop, S. Archibald, and

A. C. Staver. Reduced global fire activity due to human demography slows global warming by enhanced land carbon uptake. *Proceedings of the National Academy of Sciences*, 119(20): e2101186119, May 2022. doi: 10.1073/pnas.2101186119.

J. Yang, H. Tian, B. Tao, W. Ren, J. Kush, Y. Liu, and Y. Wang. Spatial and temporal patterns of global burned area in response to anthropogenic and environmental factors: Reconstructing global fire history for the 20th and early 21st centuries. *Journal of Geophysical Research: Biogeosciences*, 119(3):249–263, Feb. 2014. ISSN 2169-8961. doi: 10.1002/2013JG002532.

Y. Yin, A. A. Bloom, J. Worden, S. Saatchi, Y. Yang, M. Williams, J. Liu, Z. Jiang, H. Worden, K. Bowman, C. Frankenberg, and D. Schimel. Fire decline in dry tropical ecosystems enhances decadal land carbon sink. *Nat Commun*, 11(1):1900, Apr. 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-15852-2.

J.-H. Yoon, S.-Y. Wang, R. Gillies, L. Hipps, B. Kravitz, and P. J. Rasch. Extreme Fire Season in California: A Glimpse Into the Future? [in "Explaining Extremes of 2014 from a Climate Perspective"]. *Bulletin of the American Meteorological Society*, 96:S5–S9, Nov. 2015. doi: 10.1175/BAMS-D-15-00114.1.

V. Zalles, M. C. Hansen, P. V. Potapov, S. V. Stehman, A. Tyukavina, A. Pickens, X.-P. Song, B. Adusei, C. Okpa, R. Aguilar, N. John, and S. Chavez. Near doubling of Brazil's intensive row crop area since 2000. *Proceedings of the National Academy of Sciences*, 116(2):428–435, Jan. 2019. doi: 10.1073/pnas.1810301115.

J. Zeng, Z. Li, Q. Chen, H. Bi, J. Qiu, and P. Zou. Evaluation of remotely sensed and reanalysis soil moisture products over the Tibetan Plateau using in-situ observations. *Remote Sensing of Environment*, 163:91–110, June 2015. ISSN 0034-4257. doi: 10.1016/j.rse.2015.03.008.

T. Zhang, M. Wooster, M. de Jong, and W. Xu. How Well Does the 'Small Fire Boost' Methodology Used within the GFED4.1s Fire Emissions Database Represent the Timing, Location and Magnitude of Agricultural Burning? *Remote Sensing*, 10(6):823, May 2018. ISSN 2072-4292. doi: 10.3390/rs10060823.

Q. Zhao and T. Hastie. Causal Interpretations of Black-Box Models. *Journal of Business & Economic Statistics*, 0(0):1–10, June 2019. ISSN 0735-0015. doi: 10.1080/07350015.2019.1624293.

B. Zheng, P. Ciais, F. Chevallier, E. Chuvieco, Y. Chen, and H. Yang. Increasing forest fire emissions despite the decline in global burned area. *Science Advances*, 7(39):eabh2646, Sept. 2021. doi: 10.1126/sciadv.abh2646.