Geoscientific
Model Development

A non-linear Granger-causality framework to investigate climate–vegetation dynamics

Christina Papagiannopoulou[1], Diego G. Miralles[2,3], Stijn Decubber[1], Matthias Demuzere[2], Niko E. C. Verhoest[2], Wouter A. Dorigo[4], and Willem Waegeman[1]

[1]Depart. of Mathematical Modelling, Statistics and Bioinformatics, Ghent University, Ghent, Belgium
[2]Laboratory of Hydrology and Water Management, Ghent University, Ghent, Belgium
[3]Depart. of Earth Sciences, VU University Amsterdam, Amsterdam, the Netherlands
[4]Depart. of Geodesy and Geo-Information, Vienna University of Technology, Vienna, Austria

Correspondence to: Christina Papagiannopoulou (christina.papagiannopoulou@ugent.be)

Abstract. Satellite Earth observation has led to the creation of global climate data records of many important environmental and climatic variables. These come in the form of multivariate time series with different spatial and temporal resolutions. Data of this kind provide new means to further unravel the influence of climate on vegetation dynamics. However, as advocated in this article, commonly used statistical methods are often too simplistic to represent complex climate–vegetation relationships due to linearity assumptions. Therefore, as an extension of linear Granger-causality analysis, we present a novel non-linear framework consisting of several components, such as data collection from various databases, time series decomposition techniques, feature construction methods, and predictive modelling by means of random forests. Experimental results on global data sets indicate that, with this framework, it is possible to detect non-linear patterns that are much less visible with traditional Granger-causality methods. In addition, we discuss extensive experimental results that highlight the importance of considering non-linear aspects of climate–vegetation dynamics.

1 Introduction

Vegetation dynamics and the distribution of ecosystems are largely driven by the availability of light, temperature, and water; thus, they are mostly sensitive to climate conditions (Nemani et al., 2003; Seddon et al., 2016; Papagiannopoulou et al., 2017). Meanwhile, vegetation also plays a crucial role in the global climate system. Plant life alters the characteristics of the atmosphere through the transfer of water vapour, exchange of carbon dioxide, partition of surface net radiation (e.g. albedo), and impacts on wind speed and direction (Nemani et al., 2003; McPherson et al., 2007; Bonan, 2008; Seddon et al., 2016; Papagiannopoulou et al., 2017). Because of the strong two-way relationship between terrestrial vegetation and climate variability, predictions of future climate can be improved through a better understanding of the vegetation response to past climate variability.

The current wealth of Earth observation data can be used for this purpose. Nowadays, independent sensors on different platforms collect optical, thermal, microwave, altimetry, and gravimetry information, and are used to monitor vegetation, soils, oceans, and atmosphere (e.g. Su et al., 2011; Lettenmaier et al., 2015; McCabe et al., 2017). The longest composite records of environmental and climatic variables already span up to 35 years, enabling the study of multidecadal climate–biosphere interactions. Simple correlation statistics and multilinear regressions using some of these data sets have led to important steps forward in understanding the links between vegetation and climate (e.g. Nemani et al., 2003; Barichivich et al., 2014; Wu et al., 2015). However, these methods in general are insufficient when it comes to assessing causality, particularly in systems like the land–atmosphere continuum in which complex feedback mechanisms are involved. A commonly used alternative consists of Granger-causality modelling (Granger, 1969). Analyses of this kind have been applied in climate attribution studies to

investigate the influence of one climatic variable on another, e.g. the Granger-causal effect of CO_2 on global temperature (Triacca, 2005; Kodra et al., 2011; Attanasio, 2012), of vegetation and snow coverage on temperature (Kaufmann et al., 2003), of sea surface temperatures on the North Atlantic Oscillation (Mosedale et al., 2006), or of the El Niño–Southern Oscillation on the Indian monsoon (Mokhov et al., 2011). Nonetheless, Granger causality should not be interpreted as "real causality"; one assumes that a time series A Granger causes a time series B if the past of A is helpful in predicting the future of B (see Sect. 2 for a more formal definition). However, the underlying statistical model that is commonly considered in such a context is a linear vector autoregressive model, which is (again), by definition, linear; see, e.g. Shahin et al. (2014); Chapman et al. (2015).

In this article, we show new experimental evidence that advocates the need non-linear methods to study climate–vegetation dynamics due to the non-linear nature of these interactions (Foley et al., 1998; Zeng et al., 2002; Verbesselt et al., 2016). To this end, we have assembled a large, comprehensive database, comprising various global data sets of temperature, radiation, and precipitation, originating from multiple online resources. We use the Normalized Difference Vegetation Index (NDVI) to characterize vegetation, which is commonly used as a proxy of plant productivity (Myneni et al., 1997; Nemani et al., 2003). We followed an inclusive data collection approach, aiming to consider all available data sets with a worldwide coverage, and at least a 30-year time span and monthly temporal resolution (Sect. 3). Our novel non-linear Granger-causality framework is used for finding climatic drivers of vegetation and consists of several steps (Sect. 2). In a first step, we apply time series decomposition techniques to the vegetation and the various climatic time series to isolate seasonal cycles, trends, and anomalies. Subsequently, we explore various techniques for constructing more complex features from the decomposed climatic time series. In a final step, we run a Granger-causality analysis on the NDVI anomalies, while replacing traditional linear vector autoregressive models with random forests. This framework allows for modelling non-linear relationships and prevents overfitting. The results of the global application of our framework are discussed in Sect. 4.

2 A Granger-causality framework for geosciences

2.1 Linear Granger causality revisited

We start with a formal introduction to Granger causality for the case of two times series, denoted as $x = [x_1, x_2, \ldots, x_N]$ and $y = [y_1, y_2, \ldots, y_N]$, with N being the length of the time series. In this work, y alludes to the NDVI anomaly time series at a given pixel, whereas x can represent the time series of any climatic variable at that pixel (e.g. temperature, precipitation, radiation). Granger causality can be interpreted as

predictive causality, for which one attempts to forecast y_t (at the specific timestamp t) given the values of x and y in previous timestamps. Granger (1969) postulated that x causes y if the autoregressive forecast of y improves when information of x is taken into account. In order to make this definition more precise, it is important to introduce a performance measure to evaluate the forecast. Below, we will work with the coefficient of determination R^2, which is here defined as follows:

$$R^2(y, \hat{y}) = 1 - \frac{\text{RSS}}{\text{TSS}} = 1 - \frac{\sum_{i=P+1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=P+1}^{N}(y_i - \bar{y})^2}, \qquad (1)$$

where y represents the observed time series, $\bar{y}$ is the mean of this time series, $\hat{y}$ is the predicted time series obtained from a given forecasting model, and P is the length of the lag-time moving window. Therefore, the R^2 can be interpreted as the fraction of explained variance by the forecasting model, and it increases when the performance of the model increases, reaching the theoretical optimum of 1 for an error-free forecast and being negative when the predictions are less representative of the observations than the mean of the observations. Using R^2, one can now define Granger causality in a more formal way.

Definition 1. *We say that time series x Granger causes y if $R^2(y, \hat{y})$ increases when $x_{t-1}, x_{t-2}, \ldots, x_{t-P}$ are included in the prediction of y_t, in contrast to considering $y_{t-1}, y_{t-2}, \ldots, y_{t-P}$ only, where P is the lag-time moving window.*

In climate sciences, linear vector autoregressive (VAR) models are often employed to make forecasts (Stock and Watson, 2001; Triacca, 2005; Kodra et al., 2011; Attanasio, 2012). A linear VAR model of order P boils down to the following representation:

$$\begin{bmatrix} y_t \\ x_t \end{bmatrix} = \begin{bmatrix} \beta_{01} \\ \beta_{02} \end{bmatrix} + \sum_{p=1}^{P} \begin{bmatrix} \beta_{11p} & \beta_{12p} \\ \beta_{21p} & \beta_{22p} \end{bmatrix} \begin{bmatrix} y_{t-p} \\ x_{t-p} \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \end{bmatrix}, \qquad (2)$$

with β_{ij} being parameters that need to be estimated and ϵ_1 and ϵ_2 referring to two white noise error terms. This model can be used to derive the predictions required to determine Granger causality. In that sense, time series x Granger causes time series y if at least one of the parameters β_{12p} for any p significantly differs from 0. Specifically, and since we are focusing on the vegetation time series as the only target, the following two models are compared:

$$y_t = \hat{y}_t + \epsilon_1 = \beta_{01} + \sum_{p=1}^{P} \left(\beta_{11p} y_{t-p} + \beta_{12p} x_{t-p} \right) + \epsilon_1 \quad (3)$$

$$y_t = \hat{y}_t + \epsilon_1 = \beta_{01} + \sum_{p=1}^{P} \beta_{11p} y_{t-p} + \epsilon_1. \qquad (4)$$

We will refer to Eq. (3) as the "full model" and to Eq. (4) as the "baseline model", since the former incorporates all available information and the latter only information of y.

Comparing the above two models, x Granger causes y if the full model manifests a substantially better predictive performance in terms of R^2 than the baseline model. To this end, statistical tests can be employed, for which one typically assumes that the errors in the model follow a Gaussian distribution (Maddala and Lahiri, 1992). However, our above definition differs from the perspective in research papers that develop statistical tests for Granger causality (Hacker and Hatemi-J, 2006), because we intend to move away from statistical hypothesis testing, since the assumptions behind such testing are typically violated when working with climate data where neither variables nor observational techniques are fully independent from each other in most cases, and errors are not normally distributed (see Sect. 2.4 for further discussion).

In climate studies, the Granger-causal relationship between two time series x and y has often been investigated in the bivariate setting (Elsner, 2006, 2007; Kodra et al., 2011; Attanasio, 2012; Attanasio et al., 2012). However, such an analysis might lead to incorrect conclusions, because additional (confounding) effects exerted by other climatic or environmental variables are not taken into account (Geiger et al., 2015). This problem can be mitigated by considering time series of additional variables. For example, let us assume one has observed a third variable w, which might act as a confounder in deciding whether x Granger causes y. The above definition then naturally extends as follows.

Definition 2. *We say that time series x Granger causes y conditioned on time series w if $R^2(y, \hat{y})$ increases when $x_{t-1}, x_{t-2}, \ldots, x_{t-P}$ are included in the prediction of y_t, in contrast to considering $y_{t-1}, y_{t-2}, \ldots, y_{t-P}$ and $w_{t-1}, w_{t-2}, \ldots, w_{t-P}$ only, where P is the lag-time moving window.*

Similarly as above, we refer to the two models as full and baseline model, respectively. Therefore, in the trivariate setting, Granger causality might be tested using the following linear VAR model:

$$
\begin{bmatrix} y_t \\ x_t \\ w_t \end{bmatrix} = \begin{bmatrix} \beta_{01} \\ \beta_{02} \\ \beta_{03} \end{bmatrix} + \sum_{p=1}^{P} \begin{bmatrix} \beta_{11p} & \beta_{12p} & \beta_{13p} \\ \beta_{21p} & \beta_{22p} & \beta_{23p} \\ \beta_{31p} & \beta_{32p} & \beta_{33p} \end{bmatrix} \begin{bmatrix} y_{t-p} \\ x_{t-p} \\ w_{t-p} \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \epsilon_3 \end{bmatrix}, \tag{5}
$$

where a causal relationship between x and y exists if at least one β_{12p} significantly differs from 0. As previously mentioned, the time series w might also have a causal effect on y and be correlated with x. For this reason, w should be included in both models (baseline and full), so that the method can cope with cross-correlations between predictors or, in our case, between the climatic drivers of vegetation anomalies. An extension of this definition for more than three time series is straightforward.

2.2 Overfitting and out-of-sample testing

It is well known in the statistical literature that predictions made on in-sample data, i.e. the same data that were used to fit the statistical model, tend to be optimistic. This process is often referred to as overfitting; i.e. by definition, the fitting process leads to parameter values that cause the model to mimic the observed data as closely as possible (Friedman et al., 2001). In the context of Granger-causality analysis, overfitting will occur more prominently in the multivariate case, when the number of considered time series increases. The results in Sect. 4 are based on multivariate analysis; thus, they are vulnerable to overfitting; the situation further aggravates when switching from linear to non-linear models, because then the number of parameters typically increases to allow for a more flexible functional model form.

To prevent overfitting, out-of-sample data should be used in evaluating the predictive performance in Granger-causality studies (Gelper and Croux, 2007). The most straightforward procedure for creating out-of-sample data is to separate the time frame into two parts, a training set and a test set, which typically constitute the first and last halves of the time frame. A few authors have adopted this approach for climatic attribution (Attanasio et al., 2012; Pasini et al., 2012); however, satellite Earth observation time series are usually too short to allow for train-test splitting in that fashion. An alternative approach, which uses the available data in an efficient manner, is cross-validation. To this end, the time frame is divided into a number of short intervals, typically a few years of data, in which one interval serves as a test set, while all remaining data are used for parameter fitting. This procedure is repeated until all intervals have served once as a test set, and the prediction errors obtained in each round are aggregated so that one global performance measure can be computed. We direct the reader to Michaelsen (1987) and Von Storch and Zwiers (2001) for further discussion.

The inclusion of a regularization term in the fitting process of over-parameterized linear models will avoid overfitting. Typical regularizers that shrink the parameter vectors of linear models towards 0 are L2 norms (as in ridge regression), L1 norms (as in least absolute shrinkage and selection operator (LASSO) models), or a combination of the two norms (as in elastic nets) (Friedman et al., 2001). Translated to VAR models, this implies that one should impose restrictions on the parameter matrix of Eq. (5), as done in the recent theoretical paper of Gregorova et al. (2015). In this work, we want to identify causal relationships between a vegetation time series and various climatic time series. Hence, there is only one target variable of interest, and a simpler approach can be adopted. Denoting the vegetation time series by y, one can mimic in the trivariate setting a VAR model by means of three autoregressive ridge regression models:

$$y_t = \hat{y}_t + \epsilon_1 = \beta_{01}$$
$$+ \sum_{p=1}^{P} \left(\beta_{11p} y_{t-p} + \beta_{12p} x_{t-p} + \beta_{13p} w_{t-p} \right) + \epsilon_1 \qquad (6)$$

$$x_t = \hat{x}_t + \epsilon_2 = \beta_{02}$$
$$+ \sum_{p=1}^{P} \left(\beta_{21p} y_{t-p} + \beta_{22p} x_{t-p} + \beta_{23p} w_{t-p} \right) + \epsilon_2 \qquad (7)$$

$$w_t = \hat{w}_t + \epsilon_3 = \beta_{03}$$
$$+ \sum_{p=1}^{P} \left(\beta_{31p} y_{t-p} + \beta_{32p} x_{t-p} + \beta_{33p} w_{t-p} \right) + \epsilon_3. \qquad (8)$$

In this article, we aim to detect the climate drivers of vegetation and not the feedback of vegetation on climate (see, e.g. Green et al., 2017). Therefore, it suffices to retain Eq. (6) in our analysis as is stated above for the trivariate case (Eq. 5). Concatenating all parameters of this model into a vector $\boldsymbol{\beta} = [\beta_{01}, \beta_{11p}, \ldots, \beta_{13p}]$, one fits the parameters in ridge regression by solving the following optimization problem:

$$\min_{\boldsymbol{\beta}} \sum_{P+1}^{N} (y_t - \hat{y}_t)^2 + \lambda ||\boldsymbol{\beta}||^2, \qquad (9)$$

with λ being a regularization parameter, that is tuned using a validation set or nested cross-validation, and $||\boldsymbol{\beta}||^2$ being a penalty term, i.e. the squared L2 norm of the coefficient vector. The sum only starts at $P + 1$ because a moving window of P lags is considered. For simplicity, we describe the above approach for the trivariate setting, even though the total number of variables used in our study is a lot larger (see Sect. 3); nonetheless, extensions to the multivariate setting are straightforward.

2.3 Non-linear Granger causality

The methodology that we develop in this paper is closely connected to the methods explained in the previous section. However, as we hypothesize that the relationships between climate and vegetation can be highly non-linear (Foley et al., 1998; Zeng et al., 2002; Verbesselt et al., 2016), we also replace the linear VAR models in the Granger-causality framework with non-linear machine learning models. In other fields, such as neuroscience, kernel methods or other non-linear models have been used for the investigation of non-linear Granger-causality relationships between time series (Ancona et al., 2004; Marinazzo et al., 2008). In our analysis, we use simple non-linear methods that are applicable to large data sets. More sophisticated approaches typically do not scale well enough in global climate–vegetation data sets. Therefore, in our work, the machine learning algorithm we choose is random forests due to its excellent computational scalability (Breiman, 2001). Random forests is a well-known

method that has shown its merits in diverse application domains and has successfully been applied to Earth observations in both classification and regression problems (Dorigo et al., 2012; Rodriguez-Galiano et al., 2012; Loosvelt et al., 2012a, b). Briefly summarized, the random forest algorithm forms a combination of multiple decision trees, where each tree contributes a single vote to the final output, which is the most frequent class (for classification problems) or the average (for regression problems).

Compared to most application domains where random forests are applied, we employ the algorithm in a slightly different way as an autoregressive non-linear method for time series forecasting. In practice, this means that we replace the full and baseline linear model of Sect. 2.1 by a random forest model. At each pixel, the vegetation time series is still considered as a response variable, and the various climate time series serve as predictor variables (see Sect. 3.1 for an overview of our database). For a given value of the NDVI time series $\boldsymbol{y}$ at timestamp t, we investigate properties of the different predictor time series – temperature, radiation, etc. – by considering a moving window including a number of previous months (Fig. 1). In this way, the definition of Granger causality in Sect. 2.1 is adopted. Any climatic time series $\boldsymbol{x}$ Granger causes vegetation time series $\boldsymbol{y}$ if the predictive performance in terms of R^2 improves when the moving window $x_{t-1}, x_{t-2}, \ldots, x_{t-P}$ is incorporated in the random forests, in contrast to considering $y_{t-1}, y_{t-2}, \ldots, y_{t-P}$ and $w_{t-1}, w_{t-2}, \ldots, w_{t-P}$ only. Analogous to the linear case, we will speak of a full random forest model when all variables are taken into account and of a baseline random forest model when only the moving window $y_{t-1}, y_{t-2}, \ldots, y_{t-P}$ of $\boldsymbol{y}$ is considered as a predictor. In Fig. 1, this principle is extended to four time series. The baseline random forest predictions of NDVI at t_1 are based on the observations from the green moving window only, whereas the full random forest model includes the three red moving windows as well.

In our experiments, we treat each continental pixel as a separate problem and use the Scikit-learn library (Pedregosa et al., 2011) for the random forest regressor implementation, with the number of trees equal to 100 and the maximum number of predictor variables per node equal to the square root of the total number of predictor variables. Changes in these parameters or in the randomness of the algorithm do not cause substantial changes in the results (not shown). Model performance is assessed by means of 5-fold cross-validation. The window length is fixed to 12 months because initial experimental results revealed that longer time windows did not lead to improvements in the predictions (results omitted). Finally, we also experimented with techniques that exploit spatial correlations to improve the predictive performance of the model (see Sect. 4.3).
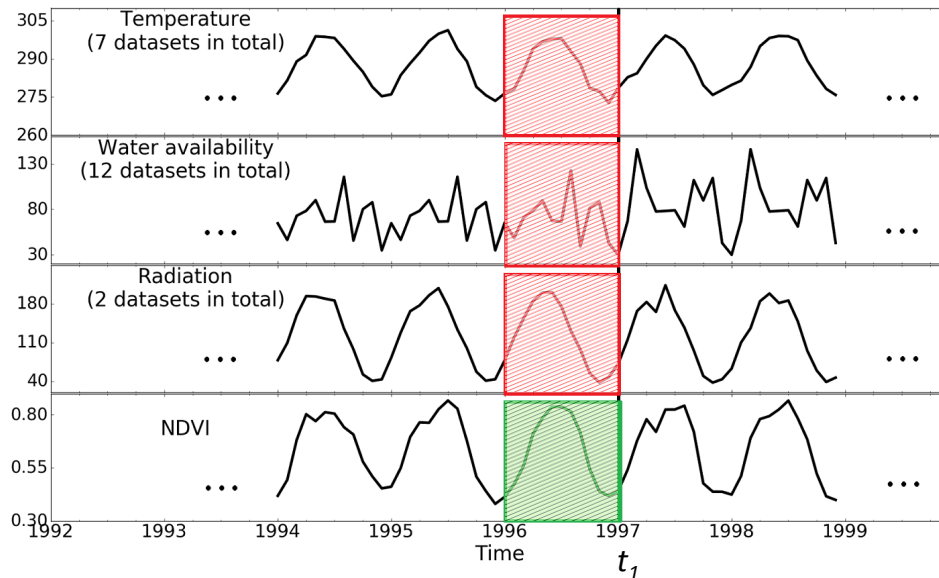
Figure 1. An illustrative example of the moving window approach considered in the analysis of vegetation drivers at a given timestamp t_1. Here, NDVI takes the role of the time series y in Eq. (3). In addition, three climate predictor time series are shown. The baseline random forest model only considers the green moving window, whereas the full random forest model includes the red moving windows as well. The pixel corresponds to a location in North America (lat: 37.5°, long: −87.5°).

2.4 Granger-causal inference

Generally, the null hypothesis (H_0) of Granger causality is that the baseline model has equal prediction error as the full model. Alternatively, if the full model predicts the target variable y significantly better than the baseline model, H_0 is rejected. In some applications, inference is drawn in VAR by testing for significance of individual model parameters. Other studies have used likelihood-ratio tests, in which the full and baseline models are nested models (Mosedale et al., 2006). However, in both cases, the models are trained and evaluated on the same in-sample data. As it has been discussed above, the performance of any Granger-causal model should be validated on out-of-sample data to avoid overfitting (see Sect. 2.2). Therefore, the null hypothesis of noncausality in the formulation stated above should be tested for by comparing out-of-sample prediction errors. To this end, statistical tests have been proposed and applied both in the econometric literature as well as in Granger-causality studies in the context of climate science. These kinds of tests, which compare out-of-sample prediction errors, are available for models for which parameter estimation is done through ordinary least squares or maximum likelihood estimation (Attanasio et al., 2013). Moreover, the asymptotic and finite-sample properties of a battery of tests for comparing forecasting accuracies of different models have been studied and, more recently, further tests aiming specifically at nested models have been proposed (Clark and McCracken, 2001). Unfortunately, all the tests mentioned above were designed to compare the out-of-sample prediction errors of lin-

ear parametric models (McCracken, 2007). In climate, relations between variables are highly non-linear and tend to become even more non-linear as the temporal resolution of the data becomes finer (Attanasio et al., 2013). Therefore, it would be convenient to have at our disposal a statistical test to assess the significance of any quantitative evidence of climate (Granger) causing vegetation anomalies. Ideally, the test would be model independent so that any non-linear model could be used. One well-known model-independent test to compare the accuracy of two forecasts is the Diebold–Mariano test (DM test) (Diebold, 2015). Although its application to Granger causality is promising, the test does not hold for nested models, because under H_0 the prediction errors from two nested models are exactly the same and perfectly correlated (McCracken, 2007). An alternative approach for comparing the predictive performance of different models is to use resampling methods such as the bootstrap or schemes such as 5×2 cross-validation (Dietterich, 1998). Methods based on the bootstrap have been used before in Granger-causality studies with climate data (Diks and Mudelsee, 2000; Attanasio et al., 2013). However, these results need to be interpreted with care because, by increasing the number of bootstrap samples, the power of any paired test (such as the Wilcoxon signed rank test) to detect significant differences between the error distributions of both models (full and baseline) increases as well. For these reasons, we conclude that developing a statistical test that is able to handle non-stationary time series and non-linear models is not a trivial task. To the best of our knowledge, no such test exists in the current literature. In this paper, we focus on express-

ing Granger causality in a quantitative instead of a qualitative way and stress the gained improvement with the use of a non-linear model.

3 Database creation and variable construction

3.1 Global data sets

Our non-linear Granger-causality framework is used to disentangle the effect of past climate variability on global vegetation dynamics. To this end, climate data sets of observational nature – mostly based on satellite and in situ observations – have been assembled to construct time series (see Sect. 3.3) that are then used to predict NDVI anomalies following the linear and non-linear causality frameworks described in Sect. 2. Data sets have been selected from the current pool of satellite and in situ observations on the basis of meeting a series of spatiotemporal requirements: (a) expected relevance of the variable for driving vegetation dynamics, (b) multidecadal record and global coverage available, and (c) adequate spatial and temporal resolution. The selected data sets can be classified into three different categories: water availability (including precipitation, snow water equivalent, and soil moisture data sets), temperature (both for the land surface and the near-surface atmosphere), and radiation (considering different radiative fluxes independently). Rather than using a single data set for each variable, we have collected all data sets meeting the above requirements. This has led to a total of 21 different data sets which are listed in Table 1. They span the study period 1981–2010 at the global scale and have been converted to a common monthly temporal resolution and $1° \times 1°$ latitude–longitude spatial resolution. To do so, we have used averages to resample original data sets found at finer native resolution and linear interpolation to resample coarser-resolution ones.

For temperature, we consider seven different products based on in situ and satellite data: Climate Research Unit (CRU-HR) (Harris et al., 2014), University of Delaware (UDel) (Willmott and Matsuura, 2001), NASA Goddard Institute for Space Studies (GISS) (Hansen et al., 2010), merged land-ocean surface temperature (MLOST) (Smith et al., 2008), International Satellite Cloud Climatology Project (ISCCP) (Rossow and Duenas, 2004), and global land surface temperature data (CFSR-Land) (Coccia et al., 2015). We also included one reanalysis data set, the European Centre for Medium-Range Weather Forecasts (ECMWF) ERA-Interim (Dee et al., 2011). In the case of precipitation, eight products have been collected. Four of them result from the merging of in situ data only: Climate Research Unit (CRU-HR) (Harris et al., 2014), University of Delaware (UDel) (Willmott and Matsuura, 2001), Climate Prediction Center Unified analysis (CPC-U) (Xie et al., 2007), and the Global Precipitation Climatology Centre (GPCC) (Schneider et al., 2011). The rest result from a combination of in situ

Table 1. Data sets used in our experiments. Basic data set characteristics are provided, including the native spatial and temporal resolutions.

Variable	Product name	Spatial resolution	Temporal resolution	Primary data source	Reference
Temperature	CRU-HR (https://crudata.uea.ac.uk/cru/data/hrg/)	0.5°	monthly	in situ	Harris et al. (2014)
	UDel (https://www.esrl.noaa.gov/psd/data/gridded/data.UDel_AirT_Precip.html)	0.5°	monthly	in situ	Willmott and Matsuura (2001)
	ISCCP (https://isccp.giss.nasa.gov/pub/data/D2Tars/)	1°	monthly	satellite	Rossow and Duenas (2004)
	ERA-Interim (http://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=sfc/)	0.75°	3-hourly	reanalysis	Dee et al. (2011)
	GISS (https://data.giss.nasa.gov/gistemp/)	2°	monthly	in situ	Hansen et al. (2010)
	MLOST (https://www.esrl.noaa.gov/psd/data/gridded/data.mlost.html)	5°	monthly	in situ	Smith et al. (2008)
	CFSR-Land (http://hydrology.princeton.edu/getdata.php?dataid=9)	0.5°	daily	satellite	Coccia et al. (2015)
Water availability	CRU-HR (https://crudata.uea.ac.uk/cru/data/hrg/)	0.5°	monthly	in situ	Harris et al. (2014)
	MSWEP (http://www.gloh2o.org/)	0.25°	3-hourly	satellite/in situ	Beck et al. (2017)
	UDel (https://www.esrl.noaa.gov/psd/data/gridded/data.UDel_AirT_Precip.html)	0.5°	monthly	in situ	Willmott and Matsuura (2001)
	CMAP (https://www.esrl.noaa.gov/psd/data/gridded/data.cmap.html)	2.5°	monthly	satellite/in situ	Xie and Arkin (1997)
	CPC-U (https://climatedataguide.ucar.edu/climate-data/cpc-unified-gauge-based-analysis-global-daily-precipitation)	0.25°	daily	in situ	Xie et al. (2007)
	GPCC (http://www.dwd.de/EN/ourservices/gpcc/gpcc.html)	0.5°	monthly	in situ	Schneider et al. (2011)
	GPCP (https://www.esrl.noaa.gov/psd/data/gridded/data.gpcp.html)	2.5°	monthly	satellite/in situ	Adler et al. (2003)
	ERA-Interim (http://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=sfc/)	0.75°	3-hourly	reanalysis	Dee et al. (2011)
	GLEAM (http://www.gleam.eu/)	0.25°	daily	satellite	Miralles et al. (2011)
	ESA CCI-PASSIVE (http://www.esa-soilmoisture-cci.org/node/145)	0.25°	daily	satellite	Dorigo et al. (2017)
	ESA CCI-COMBINED (http://www.esa-soilmoisture-cci.org/node/145)	0.25°	daily	satellite	Liu et al. (2012)
	GlobSnow (http://www.globsnow.info/index.php?page=Data)	0.25°	daily	satellite	Luojus et al. (2010)
Radiation	SRB (https://eosweb.larc.nasa.gov/project/srb/srb_table)	1°	3-hourly	satellite	Stackhouse et al. (2004)
	ERA-Interim (http://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=sfc/)	0.75°	3-hourly	reanalysis	Dee et al. (2011)
Greenness (NDVI)	GIMMS (https://ecocast.arc.nasa.gov/data/pub/gimms/)	0.25°	monthly	satellite	Tucker et al. (2005)

and satellite data, and may include reanalysis: CPC Merged Analysis of Precipitation (CMAP) (Xie and Arkin, 1997), ERA-Interim (Dee et al., 2011), Global Precipitation Climatology Project (GPCP) (Adler et al., 2003), and Multi-Source Weighted-Ensemble Precipitation (MSWEP) (Beck et al., 2017). For radiation, two different products have been collected (considering incoming short-wave/long-wave and surface net radiation as different time series): the first is the NASA Global Energy and Water cycle Exchanges (GEWEX) surface radiation budget (SRB) (Stackhouse et al., 2004) based on satellite data, and the second is the ERA-Interim reanalysis (Dee et al., 2011). For soil moisture, we use the Global Land Evaporation Amsterdam Model (GLEAM) (Miralles et al., 2011; Martens et al., 2016) and the Climate Change Initiative (CCI) product (Liu et al., 2012, 2011). Two different soil moisture products by CCI are considered: the passive microwave data set and the combined active/passive product (Dorigo et al., 2017). Moreover, snow water equivalent data come from the GlobSnow project (Luojus et al., 2010).

To conclude, as a proxy for the state and activity of vegetation, we use the third-generation (3G) Global Inventory Modeling and Mapping Studies (GIMMS) satellite-based NDVI (Tucker et al., 2005), a commonly used long-term global record of NDVI (Beck et al., 2011). We note that this data set is used to derive the response variable in our approach (seasonal NDVI anomalies; see Sect. 3.2), while all other data sets are converted to predictor variables. The length of the NDVI record (1981–2010) sets the study period to an interval of 30 years.

3.2 Anomaly decomposition

In climate studies, Granger causality has already been applied on time series of seasonal anomalies (Attanasio, 2012; Tuttle and Salvucci, 2016). The latter may be obtained in a two-step decomposition procedure by first subtracting the seasonal cycle and then the long-term trend from the raw time series. Several competing decomposition methods have been proposed in the literature, including additive models, multiplicative models, and more sophisticated methods based on break points (see, e.g. Cleveland et al., 1990; Grieser et al., 2002; Verbesselt et al., 2010). In our framework, we used the following approach: in a first step, at each given pixel, the "raw" time series of the target variable y_t and the climate predictors $(x_t, w_t, \dots)$ are detrended linearly based on a simple linear regression with the timestamp t as a predictor variable applied to the entire study period. For the case of the target variable, this can be denoted as follows:

$$y_t \approx y_t^{Tr} = \alpha_0 + \alpha_1 t, \tag{10}$$

with α_0 and α_1 being the intersect and the slope of the linear regression, respectively. We obtain in this way the detrended time series $y_t^D = y_t - y_t^{Tr}$. This detrending is needed to remove non-stationary signals in climatic time series, and al-
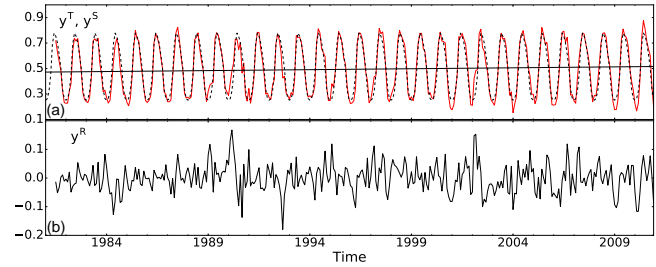

Figure 2. The three components of the NDVI time series decomposition of a specific pixel of the Northern Hemisphere (lat: 53.5°, long: 26.5°). On top are the linear trend (black continuous line) and the seasonal cycle (dashed black line) fitted on the raw data (red). On the bottom are the remaining anomalies; see text for details.

lows us to draw the emphasis to the shorter-term multi-month dynamics. By detrending, one can assure that the mean of the probability distribution does not change over time; however, other moments of the probability distribution, such as the variance, might still be time dependent. As classical statistical procedures for testing Granger causality (i.e. autoregressive model, statistical tests) are developed for stationary time series, those methods are in fact not applicable to non-stationary climate data. In a second step, after subtracting the trend from the raw time series, the seasonal cycle y_t^S is calculated. When the assumption is made that the seasonal cycle is annual and constant over time, one can simply estimate it as the monthly expectation. To this end, the multi-year average for each of the 12 months of the year is calculated. Finally, the anomalies y_t^R can then be computed by subtracting the corresponding monthly expectation from the detrended time series: $y_t^R = y_t^D - y_t^S$. This procedure is schematically represented in Fig. 2.

3.3 Predictor variable construction

We do not limit our approach to considering raw and anomaly time series of the data sets in Table 1 as predictors but also take into consideration different lag times, past cumulative values, and extreme indices (see following text). These additional predictors, here referred to as "higher-level variables", are calculated based on raw and anomaly time series. Our application of Granger causality can be interpreted as a way to identify patterns in climate during past moving windows (see Fig. 1) that are predictive with respect to the anomalies of vegetation time series. Therefore, by feeding predictor variables from previous timestamps to a linear (or non-linear) predictive model, one can identify subsequences of interest in the moving window specified for timestamp t, a technique that is similar to so-called shapelets (Ye and Keogh, 2009). In addition, vegetation dynamics may not necessarily reflect the climatic conditions from, e.g. 3 months ago, but the average of the, e.g. three antecedent months. This integrated response to antecedent environmental and cli-
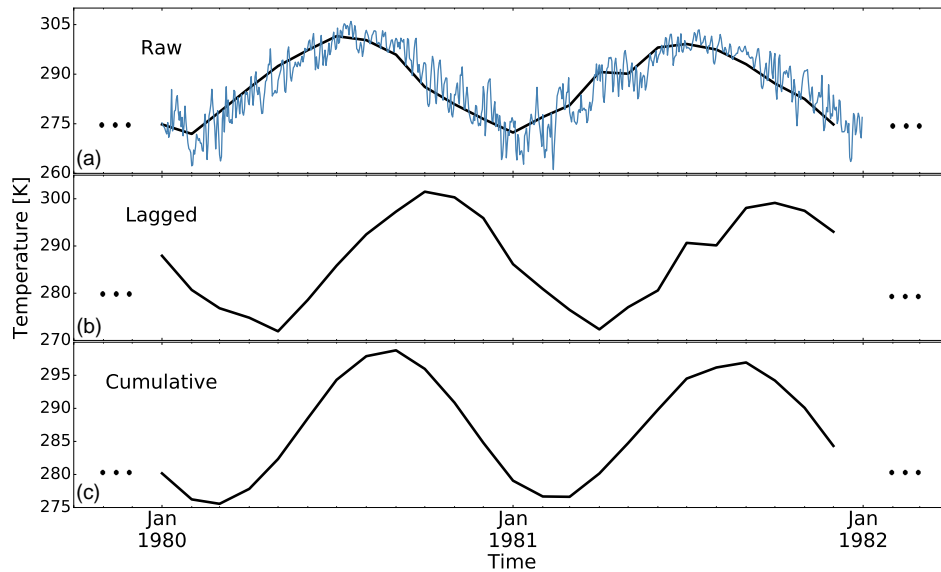
Figure 3. Example of lagged and cumulative variables extracted from a temperature time series. On top is part of a raw daily time series with its monthly aggregation. In the middle is the 4-month lag-time monthly time series. On the bottom is the corresponding 4-month cumulative variable. The pixel corresponds to a location in Kentucky, USA (lat: 37.5°, long: −87.5°).

matic conditions is referred to here as a "cumulative" response. More formally, we construct a cumulative variable of k months as the sum of time series observations in the last k months:

$$\mathrm{Cumul}[x_{t-1}, x_{t-2}, \ldots, x_{t-k}] = \sum_{p=1}^{k} x_{t-p}. \qquad (11)$$

Note that, unlike in the case of lagged variables, cumulative variables always include the period up to time t. Figure 3 illustrates an example of a 4-month cumulative variable. In our analysis, we experimented with time lags covering a wide range of time-lag values and concluded that including lags of more than 6 months did not yield substantial predictive power.

Another type of higher-level predictor variable that can be constructed from the data sets in Table 1 are extreme indices. Over the last few years, several research studies have focused on defining and indexing climate extremes (Nicholls and Alexander, 2007; Zwiers et al., 2013). As an example, the Expert Team on Climate Change Detection and Indices (ETCCDI) recommends the use of a range of extreme indices related to temperature and precipitation (Zhang et al., 2011; Donat et al., 2013). Here, we calculate a variety of analogous indices for the whole set of the collected climatic variables, based on both the raw data sets as well as on the seasonal anomalies (see Table 2). In addition, we derived lagged and cumulative predictor variables from these extremes' indices to incorporate the potential impact of climatic extremes occurring, e.g. 3 months ago, or during the previous, e.g. 3 months, respectively. All these resulting time series appear

as additional predictor variables in our non-linear Granger-causality framework (see Sect. 2.3).

Combining the different climate and environmental predictor variables described above, we obtain a database of 4571 predictor variables per 1° pixel, covering 30 years at a monthly temporal resolution.

4 Results and discussion

4.1 Detecting linear Granger-causal relationships

In a first experiment, we evaluate the extent to which climate variability Granger causes the anomalies in vegetation using a standard Granger-causality approach, in which only linear relationships between climate (predictors) and vegetation (target variable) are considered. To this end, ridge regression is used as a linear VAR model in the Granger-causality approach (note that this ridge regression will be substituted by the non-linear random forest approach in Sect. 4.2). In the application of the ridge regression, we use all climatic and environmental predictor variables (Sect. 3.3) and adopt a nested 5-fold cross-validation to properly tune the hyper parameter λ (see Eq. 9). Figure 4a shows the predictive performance of the full ridge regression model. While the model explains more than 40 % of the variability in NDVI anomalies in some regions ($R^2 > 0.4$), this is by itself not necessarily indicative of climate Granger causing the vegetation anomalies, as it may reflect simple correlations. In order to test the latter, we compare the results of the full model to a baseline model, i.e. an autoregressive ridge regression model that only uses previous values of NDVI to predict the NDVI at time t (see

Sect. 2.1). If climate Granger caused the variability of NDVI at a given pixel, the full ridge regression model (Fig. 4a) would show an increase in the predictive power over the predictions based on the baseline ridge regression model. However, the results unequivocally show that – when only linear relationships between vegetation and climate are considered – the areas for which vegetation anomalies are Granger caused by climate are very limited, involving mainly semiarid regions and central Europe (Fig. 4b).

For further comparison, we analyse the predictive performance obtained when (linear) Pearson correlation coefficients are calculated on the training data sets, selecting the highest correlation to the target variable for any of the 4571 predictor variables at each pixel. Figure 4c shows that the explained variance is again rather low and, for most regions, substantially lower than the R^2 of the baseline ridge regression model, here considered as the minimum to interpret this predictive power as Granger causal. These results indicate that, despite being routinely used as a standard tool in climate–biosphere studies (see, e.g. Nemani et al., 2003), univariate correlation analyses are unable to extract the nuances of the relationships between climate and vegetation dynamics.

4.2 Linear versus non-linear Granger causality

To analyse the effect of climate on vegetation more thoroughly, we substitute the linear ridge regression model (VAR) by the non-linear random forest model. Results in Fig. 5 highlight the differences. Compared to the results in Sect. 4.1, the predictive power substantially increases by considering non-linear relationships between vegetation and climate (Fig. 5a). This is the case for most land regions but is especially remarkable in semiarid regions of Australia, Africa, and Central and North America, which are frequently exposed to water limitations. In those regions, more that 40 % of the variance of NDVI anomalies can be explained by antecedent climate variability. These results are further investigated by Papagiannopoulou et al. (2017), who highlight the crucial role of water supply for the anomalies in vegetation greenness in these and other regions. On the other hand, the variance of NDVI explained in other areas, such as the Eurasian taiga, tropical rainforests, or China, is again below 10 %. We hypothesize two potential reasons: (a) the uncertainty in the observations used as target and predictors are typically larger in these regions (especially in tropical forests and at higher latitudes), and (b) these are regions in which vegetation anomalies are not necessarily primarily controlled by climate but may be predominantly driven by phenological and biotic factors (Hutyra et al., 2007), occurrence of wildfires (Van der Werf et al., 2010), limitations imposed by the availability of soil nutrients (Fisher et al., 2012), or agricultural practices (Liu et al., 2015). Nonetheless, the explained variance shown in Fig. 5a is again not necessarily indicative of Granger causality. As we did in Fig. 4b, in order to test whether the climatic and environmental controls do, in fact, Granger cause the vegetation anomalies, we compare the results of our full random forest model to a baseline random forest model which only uses previous values of NDVI to predict the NDVI at time t. As seen in Fig. 5b, in this case, the improvement over the baseline is unambiguous. One can conclude that – while not considering all potential control variables in our analysis – climate dynamics indeed (Granger) cause vegetation anomalies in most of the continental land surface, with a larger impact on subtropical regions and midlatitudes. Moreover, a comparison between Figs. 4b and 5b unveils that these causal relationships are highly non-linear, as expected given the distinct resistance and resilience of different ecosystems, which are reflected by a progressive response and recovery of vegetation to these perturbations (Foley et al., 1998; Zeng et al., 2002; Verbesselt et al., 2016).

For a better understanding of the results obtained by the two models, we average the performance of each model regionally. More specifically, we use the International Geosphere-Biosphere Program (IGBP) (Loveland and Belward, 1997) land cover classification to stratify the mean and variance of R^2 for both the baseline and the full model in Fig. 5 per IGBP land cover class. The bar plot in Fig. 6 shows that the full model outperforms the baseline model in all IGBP land cover classes, i.e. that Granger causality exists for all these biomes. In the parentheses, we note the number of pixels per region. The error bars indicate that the variances of the two models are analogous; i.e. they are low or high in both models in the same land cover class. For the Closed Shrublands region, one can observe the highest difference between the two models, yet only 19 pixels belong to this biome type. In savannah regions, the performance of the full model is high in comparison with other regions (see Fig. 5). On the other hand, the lowest performance improvement of the full model with respect to the baseline is observed for the regions of Deciduous Needleleaf Forests and Evergreen Broadleaf Forests. This shows that for these two regions climate is not identified as a major control over vegetation dynamics (see discussion in previous paragraph about tropical and boreal regions).

4.3 Spatial and temporal aspects

Environmental dynamics reveal their effect on vegetation at different timescales. Since the adaptation of vegetation to environmental changes requires some time, and because soil and atmosphere have a memory, a necessary aspect to investigate is the potential lag-time response of vegetation to climate dynamics which relates to the ecosystem resistance and resilience properties. The idea of exploring lag times was introduced by several studies in the past (see, e.g. Davis, 1984; Braswell et al., 1997), and it has been adopted in various studies more recently (Anderson et al., 2010; Kuzyakov and Gavrichkova, 2010; Chen et al., 2014; Rammig et al.,

Table 2. Extreme indices considered as predictive variables. These indices are derived from the raw (daily) data and the (daily) anomalies of the data sets in Table 1. We also calculate the lagged and cumulative variables from these extreme indices.

Name	Description
SD	Standard deviation of daily values per month
DIR	Difference between max and min daily value per month
Xx	Max daily value per month
Xn	Min daily value per month
Max5day	Max over 5 consecutive days per month
Min5day	Min over 5 consecutive days per month
X99p/X95p/X90p	Number of days per month over 99th/95th/90th percentile
X1p/X5p/X10p	Number of days per month under 1st/5th/10th percentile
T25C[1]	Number of days per month over 25 °C
T0C[1]	Number of days per month below 0 °C
R10mm/R20mm[2]	Number of days per month over 10/20 mm
CHD (Consecutive high-value days)	Number of consecutive days per month over 90th percentile
CLD (Consecutive low-value days)	Number of consecutive days per month below 10th percentile
CDD (Consecutive dry days)[2]	Number of consecutive days per month when precipitation < 1 mm
CWD (Consecutive wet days)[2]	Number of consecutive days per month when precipitation ≥ 1 mm
Spatial heterogeneity[3]	Difference between max and min values within 1° box

[1] Only for temperature data sets. [2] Only for precipitation data sets. [3] Only for data sets with native spatial resolution < 1° lat–long.
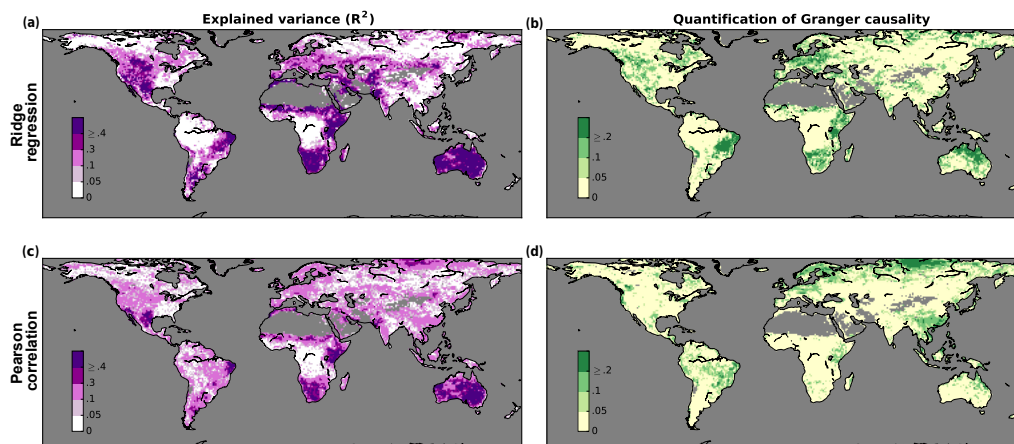

Figure 4. Linear Granger causality of climate on vegetation. **(a)** Explained variance (R^2) of NDVI anomalies based on a full ridge regression model in which all climatic variables are included as predictors. **(b)** Improvement in terms of R^2 by the full ridge regression model with respect to the baseline ridge regression model that uses only past values of NDVI anomalies as predictors; positive values indicate (linear) Granger causality. **(c)** A filter approach in which the variable with the highest squared Pearson correlation against the NDVI anomalies is selected. **(d)** Improvement in terms of R^2 by the filter approach with respect to the same baseline ridge regression model that uses only past values of NDVI anomalies.

2014). These studies indicate that lag times depend on both the specific climatic control variable and the characteristics of the ecosystem. As explained in Sect. 3.3, in our analysis shown in Figs. 4 and 5, we moved beyond traditional cross-correlations and incorporated higher-level variables in the form of cumulative and lagged responses to extreme climate. As mentioned in Sect. 3.3, our experiments indicated that lags of more than 6 months do not add extra predictive power (not shown), even though the effect of anoma-

lies in water availability on vegetation can extend for several months (Papagiannopoulou et al., 2017).

To disentangle the response of vegetation to past cumulative climate anomalies and climatic extremes, Fig. 7a visualizes the predictive performance when cumulative variables and extreme indices are not included as predictive variables in the random forest model. As shown in Fig. 7b, in almost all regions of the world the predictive performance decreases substantially compared to the full random forest model approach, i.e. using the full repository of predictors (Fig. 5a),
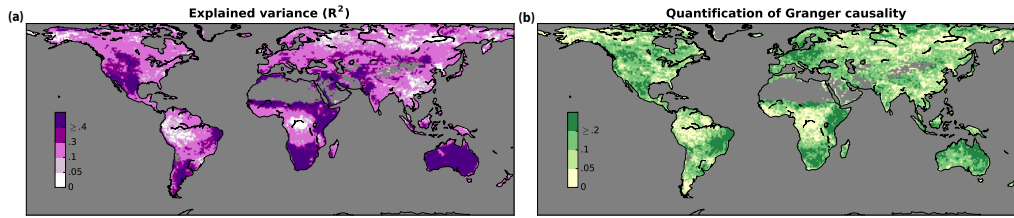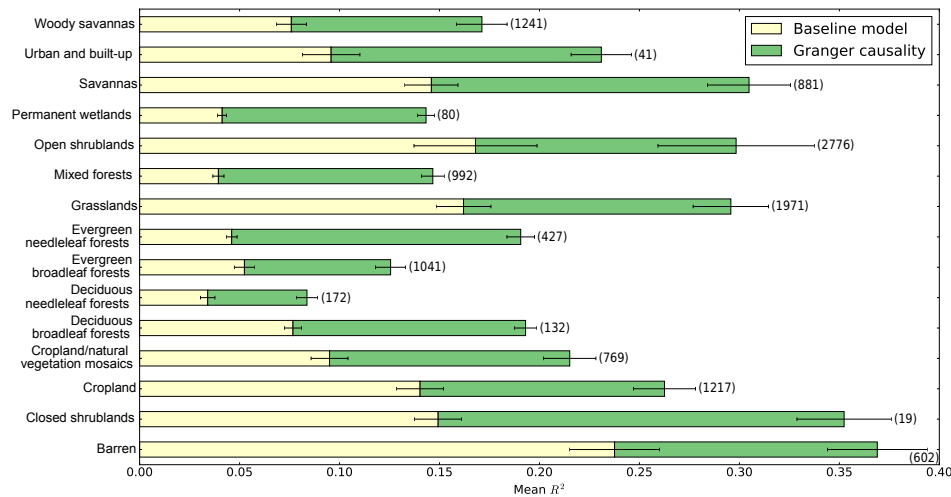
Figure 5. Non-linear Granger causality of climate on vegetation. **(a)** Explained variance (R^2) of NDVI anomalies based on a full random forest model in which all climatic variables are included as predictors. **(b)** Improvement in terms of R^2 by the full random forest model with respect to the baseline random forest model that uses only past values of NDVI anomalies as predictors; positive values indicate (non-linear) Granger causality.

Figure 6. Mean R^2 and variance per IGBP land cover class for both the baseline and full random forest model. The green part indicates the improvement in performance of the full model with respect to the baseline, i.e. the quantification of Granger causality (as in Fig. 5b). The number of pixels per IGBP class is noted in the parentheses.

especially in regions such as the Sahel, the Horn of Africa, or North America. In those regions, 10–20 % of the variability in NDVI is explained by the occurrence of prolonged anomalies and/or extremes in climate, illustrating again the non-linear responses of vegetation. For more detailed results about lagged vegetation responses for specific climate drivers and the effect of climate extremes on vegetation, the reader is referred to Papagiannopoulou et al. (2017).

Because of uncertainties in the observational records used in our study to represent climate and predict vegetation dynamics, and given that ecosystems and regional climate conditions usually extend over areas that exceed the spatial resolution of these records, one may expect that the predictive performance of our models becomes more robust when including climate information from neighbouring pixels. In addition, it is quite likely that neighbouring areas have similar climatic conditions which, in turn, affect vegetation dynamics in a similar manner. We therefore also consider an extension of our framework to exploit spatial autocorrelations, inspired by Lozano et al. (2009), who achieved spatial smoothness via an additional penalty term that punishes dissimilar-

ity between coefficients for spatial neighbours. In our analysis, we incorporate spatial autocorrelations at a given pixel by extending the predictor variables of our models with the predictor variables of the eight neighbouring pixels. We provide such an extension both for the full and the baseline random forest model. As such, for the full random forest model, a vector of 41 139 (4571 × 9) predictor variables is formed for each pixel.

Figure 7c illustrates the performance of the full random forest model that includes the spatial information. As one can observe in Fig. 7d, the explained variance of NDVI anomalies remains similar to the original model that depicts the same approach without spatial autocorrelation (Fig. 5a). While in most areas the performance slightly increases, the explained variance never improves by more than 10 %; as a result, incorporating spatial autocorrelations in our framework does not seem to further improve the quantification of Granger causality and is not considered in further applications of the framework (see Papagiannopoulou et al., 2017). A possible explanation for this result is that the model without the spatial information cannot be outperformed because
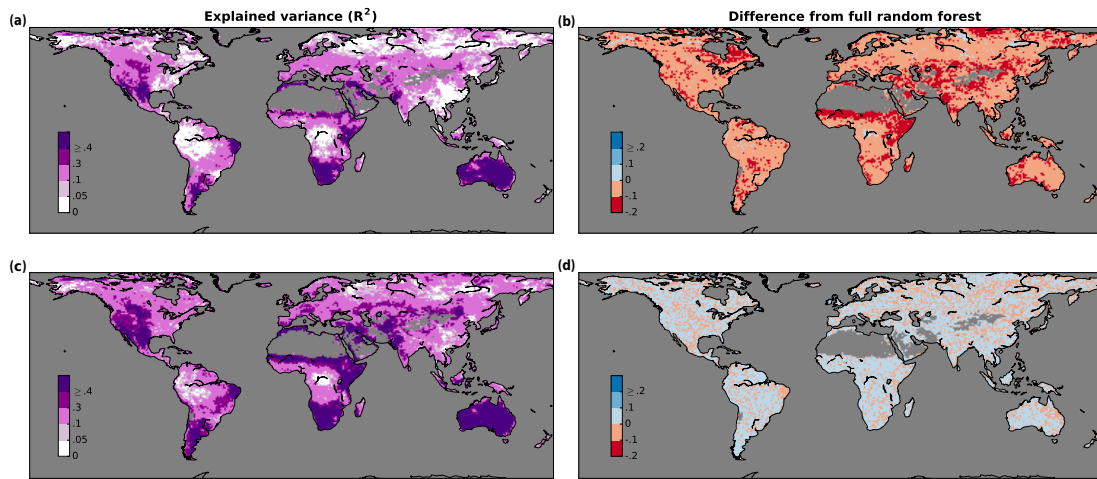
Figure 7. Analysis of spatiotemporal aspects of our framework. **(a)** Explained variance (R^2) of NDVI anomalies based on a full random forest model in which all climatic variables are included as predictors as in Fig. 5a, except for the cumulative variables and the extreme indices (see Sect. 3.3). **(b)** Difference in terms of R^2 between the model without cumulative and extreme predictors and the full random forest model in Fig. 5a. **(c)** Explained variance (R^2) of NDVI anomalies based on a full random forest model in which all climatic variables are included as predictors as in Fig. 5a, as well as the predictors from the eight nearest neighbours. **(d)** Difference in terms of R^2 between this full random forest model which includes spatial information from neighbouring pixels and the full random forest model in Fig. 5a.

Figure 8. Comparison of model performance with R^2 as the metric with the raw NDVI time series as target variable. **(a)** Full random forest model. **(b)** Improvement in terms of R^2 of the full random forest model over the baseline random forest model.

of the large dimensionality of the feature space, which may include redundant information, in combination with the low number of observations per pixel (Fig. 5a). Note that in this case the number of observations per pixel remains the same as in the original model (360 observations) while the number of predictor variables is 9 times larger.

4.4 The importance of focusing on vegetation anomalies

In Sect. 3.2, we advocated that Granger-causality analysis should target NDVI anomalies, as opposed to raw NDVI values. There are several fundamental reasons for this. First, by applying a decomposition, one can subtract long-term trends from the NDVI time series, making the resulting time series more stationary. This is absolutely needed, as existing Granger-causality tests cannot be applied for non-stationary time series. Secondly, by subtracting the seasonal cycle from the time series, one is not only able to remove a confounding factor that may contribute predictive power without bear-

ing causality but also able to remove a clear autoregressive component that can be well explained from the NDVI time series themselves. As vegetation has a strong seasonal cycle, it is not difficult to predict subsequent vegetation conditions by using the past observations of the seasonal cycle only. To corroborate this aspect, we repeat our analysis in Sect. 4.2, but this time the raw NDVI time series instead of the NDVI anomalies are considered as the target variable. We again compare the full and the baseline random forest models.

The results are visualized in Fig. 8a. As it can be observed, worldwide the R^2 is close to the optimum of 1. However, due to the overwhelming domination of the seasonal cycle, it becomes very difficult, or even impossible, to unravel any potential Granger-causal relationships with climate time series in the Northern Hemisphere; see Fig. 8b. The predictability of NDVI based on the seasonal NDVI cycle itself is already so high that nothing can be gained by adding additional climatic predictor variables (see also the large amplitude of the seasonal cycle of NDVI at those latitudes compared to the

NDVI anomalies, as illustrated in Fig. 2). Therefore, a non-linear baseline autoregressive model is able to explain most of the variance in the time series. Moreover, as observed in Fig. 1, temperature and radiation also manifest strong seasonal cycles that often coincide with the NDVI cycle. For most regions on Earth, such a stationary seasonal cycle is less present for variables such as precipitation. This can potentially yield wrong conclusions, such as that temperature in the Northern Hemisphere is driving most NDVI variability, since the two seasonal cycles have the same pattern. However, based on the above discussion, it becomes clear that results of that kind should be treated with caution: for climate data, a Granger-causality analysis should be applied after decomposing time series into seasonal anomalies.

5 Conclusions

In this paper, we introduced a novel framework for studying Granger causality in climate–vegetation dynamics. We compiled a global database of observational records spanning a 30-year time frame, containing satellite, in situ, and reanalysis-based data sets. Our approach consists of the combination of data fusion, feature construction, and non-linear predictive modelling. The choice of random forest as a non-linear algorithm has been motivated by its excellent computational scalability with regards to extremely large data sets, but could be easily replaced by any other non-linear machine learning technique, such as neural networks or kernel methods.

Our results highlight the non-linear nature of climate–vegetation interactions and the need to move beyond the traditional application of Granger causality within a linear framework. Comparisons to linear Granger-causality-based approaches indicate that the random forest framework can predict 14 % more variability of vegetation anomalies on average globally. The predictive power of the model is especially high in water-limited regions where a large part of the vegetation dynamics responds to the occurrence of antecedent rainfall. Moreover, our results indicate the need to consider multi-month antecedent periods to capture the effect of climate on vegetation, in particular to account for the effects of climate extremes on vegetation resilience. The reader is referred to Papagiannopoulou et al. (2017) for a detailed analysis of the effect of different climate predictors on the variability of global vegetation using the mathematical approach described here.

References

Adler, R. F., Huffman, G. J., Chang, A., Ferraro, R., Xie, P.-P., Janowiak, J., Rudolf, B., Schneider, U., Curtis, S., Bolvin, D., et al.: The version-2 global precipitation climatology project (GPCP) monthly precipitation analysis (1979-present), J. Hydrometeorol., 4, 1147–1167, 2003.

Ancona, N., Marinazzo, D., and Stramaglia, S.: Radial basis function approach to nonlinear Granger causality of time series, Phys. Rev. E, 70, 056221, doi:10.1103/PhysRevE.70.056221, 2004.

Anderson, L. O., Malhi, Y., Aragão, L. E., Ladle, R., Arai, E., Barbier, N., and Phillips, O.: Remote sensing detection of droughts in Amazonian forest canopies, New Phytol., 187, 733–750, 2010.

Attanasio, A.: Testing for linear Granger causality from natural/anthropogenic forcings to global temperature anomalies, Theor. Appl. Climatol., 110, 281–289, 2012.

Attanasio, A., Pasini, A., and Triacca, U.: A contribution to attribution of recent global warming by out-of-sample Granger causality analysis, Atmos. Sci. Lett., 13, 67–72, 2012.

Attanasio, A., Pasini, A., and Triacca, U.: Granger causality analyses for climatic attribution, Atmos. Clim. Sci., 3, doi:10.4236/acs.2013.34054, 2013.

Barichivich, J., Briffa, K. R., Myneni, R., van der Schrier, G., Dorigo, W., Tucker, C. J., Osborn, T. J., and Melvin, T. M.: Temperature and Snow-Mediated Moisture Controls of Summer Photosynthetic Activity in Northern Terrestrial Ecosystems between 1982 and 2011, Remote Sens., 6, 1390, doi:10.3390/rs6021390, 2014.

Beck, H. E., McVicar, T. R., van Dijk, A. I., Schellekens, J., de Jeu, R. A., and Bruijnzeel, L. A.: Global evaluation of four AVHRR–NDVI data sets: Intercomparison and assessment against Landsat imagery, Remote Sens. Environ., 115, 2547–2563, 2011.

Beck, H. E., van Dijk, A. I. J. M., Levizzani, V., Schellekens, J., Miralles, D. G., Martens, B., and de Roo, A.: MSWEP: 3-hourly 0.25° global gridded precipitation (1979–2015) by merging gauge, satellite, and reanalysis data, Hydrol. Earth Syst. Sci., 21, 589–615, doi:10.5194/hess-21-589-2017, 2017.

Bonan, G. B.: Forests and climate change: forcings, feedbacks, and the climate benefits of forests, Science, 320, 1444–1449, 2008.

Braswell, B., Schimel, D., Linder, E., and Moore, B.: The response of global terrestrial ecosystems to interannual temperature variability, Science, 278, 870–873, 1997.

Breiman, L.: Random forests, Mach. Learn., 45, 5–32, 2001.

Chapman, D., Cane, M. A., Henderson, N., Lee, D. E., and Chen, C.: A vector autoregressive ENSO prediction model, J. Climate, 28, 8511–8520, 2015.

Chen, T., De Jeu, R., Liu, Y., Van der Werf, G., and Dolman, A.: Using satellite based soil moisture to quantify the water driven variability in NDVI: A case study over mainland Australia, Remote Sens. Environ., 140, 330–338, 2014.

Clark, T. E. and McCracken, M. W.: Tests of equal forecast accuracy and encompassing for nested models, J. Econometrics, 105, 85–110, 2001.

Cleveland, R. B., Cleveland, W. S., McRae, J. E., and Terpenning, I.: STL: A seasonal-trend decomposition procedure based on loess, Journal of Official Statistics, 6, 3–73, 1990.

Coccia, G., Siemann, A. L., Pan, M., and Wood, E. F.: Creating consistent datasets by combining remotely-sensed data and land surface model estimates through Bayesian uncertainty post-processing: The case of Land Surface Temperature from {HIRS}, Remote Sens. Environ., 170, 290–305, doi:10.1016/j.rse.2015.09.010, 2015.

Davis, M. B.: Climatic instability, time, lags, and community disequilibrium, Harper & Row, 1984.

Dee, D., Uppala, S., Simmons, A., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M., Balsamo, G., Bauer, P., et al.: The ERA-Interim reanalysis: Configuration and performance of the data assimilation system, Q. J. Roy. Meteor. Soc., 137, 553–597, 2011.

Diebold, F. X.: Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of Diebold–Mariano tests, J. Bus. Econ. Stat., 33, doi:10.1080/07350015.2014.983236, 2015.

Dietterich, T. G.: Approximate statistical tests for comparing supervised classification learning algorithms, Neural Comput., 10, 1895–1923, 1998.

Diks, C. and Mudelsee, M.: Redundancies in the Earth's climatological time series, Phys. Lett. A, 275, 407–414, 2000.

Donat, M., Alexander, L., Yang, H., Durre, I., Vose, R., Dunn, R., Willett, K., Aguilar, E., Brunet, M., Caesar, J., et al.: Updated analyses of temperature and precipitation extreme indices since the beginning of the twentieth century: The HadEX2 dataset, J. Geophys. Res.-Atmos., 118, 2098–2118, 2013.

Dorigo, W., Wagner, W., Albergel, C., Albrecht, F., Balsamo, G., Brocca, L., Chung, D., Ertl, M., Forkel, M. Gruber, A., Haas, E., Hamer, P., Hirschi, M., Ikonen, J., de Jeu, R., Kidd, R., Lahoz, W., Liu, Y. Y., Miralles, D. G., Mistelbauer, T., Nicolai-Shaw, N., Parinussa, R., Pratola, C., Reimer, C., van der Schalie, R., Seneviratne, S. I., Smolander, T., and Lecomte, P.: ESA CCI Soil Moisture for improved Earth System understanding: state-of-the-art and future directions, Remote Sens. Environ., in review, 2017.

Dorigo, W., Lucieer, A., Podobnikar, T., and Čarni, A.: Mapping invasive Fallopia japonica by combined spectral, spatial, and temporal analysis of digital orthophotos, Int. J. Appl. Earth Obs., 19, 185–195, 2012.

Elsner, J. B.: Evidence in support of the climate change–Atlantic hurricane hypothesis, Geophys. Res. Lett., 33, doi:10.1029/2006GL026869, 2006.

Elsner, J. B.: Granger causality and Atlantic hurricanes, Tellus A, 59, 476–485, 2007.

Fisher, J. B., Badgley, G., and Blyth, E.: Global nutrient limitation in terrestrial vegetation, Global Biogeochem. Cy., 26, doi:10.1029/2011GB004252, 2012.

Foley, J. A., Levis, S., Prentice, I. C., Pollard, D., and Thompson, S. L.: Coupling dynamic models of climate and vegetation, Glob. Change Biol., 4, 561–579, 1998.

Friedman, J., Hastie, T., and Tibshirani, R.: The elements of statistical learning, vol. 1, Springer series in statistics Springer, Berlin, 2001.

Geiger, P., Zhang, K., Gong, M., Janzing, D., and Schölkopf, B.: Causal inference by identification of vector autoregressive processes with hidden components, in: Proceedings of 32th International Conference on Machine Learning (ICML 2015), 2015.

Gelper, S. and Croux, C.: Multivariate out-of-sample tests for Granger causality, Comput. Stat. Data An., 51, 3319–3329, 2007.

Granger, C. W.: Investigating causal relations by econometric models and cross-spectral methods, Econometrica, 37, 424–438, 1969.

Green, J. K. and et al.: Hotspots of terrestrial biosphere-atmosphere feedbacks, Nat. Geosci., doi:10.1038/NGEO2957, 2017.

Gregorova, M., Kalousis, A., Marchand-Maillet, S., and Wang, J.: Learning vector autoregressive models with focalised Granger-causality graphs, arXiv preprint https://arxiv.org/abs/1507.01978v1, 2015.

Grieser, J., Trömel, S., and Schönwiese, C.-D.: Statistical time series decomposition into significant components and application to European temperature, Theor. Appl. Climatol., 71, 171–183, 2002.

Hacker, R. S. and Hatemi-J, A.: Tests for causality between integrated variables using asymptotic and bootstrap distributions: theory and application, Appl. Econ., 38, 1489–1500, doi:10.1080/00036840500405763, 2006.

Hansen, J., Ruedy, R., Sato, M., and Lo, K.: Global surface temperature change, Rev. Geophys., 48, doi:10.1029/2010RG000345, 2010.

Harris, I., Jones, P., Osborn, T., and Lister, D.: Updated high-resolution grids of monthly climatic observations–the CRU TS3. 10 Dataset, Int. J. Climatol., 34, 623–642, 2014.

Hutyra, L. R., Munger, J. W., Saleska, S. R., Gottlieb, E., Daube, B. C., Dunn, A. L., Amaral, D. F., De Camargo, P. B., and Wofsy, S. C.: Seasonal controls on the exchange of carbon and water in an Amazonian rain forest, J. Geophys. Res.-Biogeo., 112, doi:10.1029/2006JG000365, 2007.

Kaufmann, R., Zhou, L., Myneni, R., Tucker, C., Slayback, D., Shabanov, N., and Pinzon, J.: The effect of vegetation on surface temperature: A statistical analysis of NDVI and climate data, Geophys. Res. Lett., 30, doi:10.1029/2003GL018251, 2003.

Kodra, E., Chatterjee, S., and Ganguly, A. R.: Exploring Granger causality between global average observed time series of carbon

dioxide and temperature, Theor. Appl. Climatol., 104, 325–335, 2011.

Kuzyakov, Y. and Gavrichkova, O.: REVIEW: Time lag between photosynthesis and carbon dioxide efflux from soil: a review of mechanisms and controls, Glob. Change Biol., 16, 3386–3406, 2010.

Lettenmaier, D. P., Alsdorf, D., Dozier, J., Huffman, G. J., Pan, M., and Wood, E. F.: Inroads of remote sensing into hydrologic science during the WRR era, Water Resour. Res., 51, 7309–7342, doi:10.1002/2015WR017616, 2015.

Liu, Y. Y., Parinussa, R. M., Dorigo, W. A., De Jeu, R. A. M., Wagner, W., van Dijk, A. I. J. M., McCabe, M. F., and Evans, J. P.: Developing an improved soil moisture dataset by blending passive and active microwave satellite-based retrievals, Hydrol. Earth Syst. Sci., 15, 425–436, doi:10.5194/hess-15-425-2011, 2011.

Liu, Y., Dorigo, W., Parinussa, R., De Jeu, R., Wagner, W., McCabe, M., Evans, J., and Van Dijk, A.: Trend-preserving blending of passive and active microwave soil moisture retrievals, Remote Sens. Environ., 123, 280–297, 2012.

Liu, Y., Pan, Z., Zhuang, Q., Miralles, D. G., Teuling, A. J., Zhang, T., An, P., Dong, Z., Zhang, J., He, D., et al.: Agriculture intensifies soil moisture decline in Northern China, Scientific reports, 5, 2015.

Loosvelt, L., Peters, J., Skriver, H., De Baets, B., and Verhoest, N. E.: Impact of reducing polarimetric SAR input on the uncertainty of crop classifications based on the random forests algorithm, IEEE T. Geosci. Remote Sens., 50, 4185–4200, 2012a.

Loosvelt, L., Peters, J., Skriver, H., Lievens, H., Van Coillie, F. M., De Baets, B., and Verhoest, N. E. C.: Random Forests as a tool for estimating uncertainty at pixel-level in SAR image classification, Int. J. Appl. Earth Obs., 19, 173–184, 2012b.

Loveland, T. and Belward, A.: The IGBP-DIS global 1km land cover data set, DISCover: first results, Int. J. Remote Sens., 18, 3289–3295, 1997.

Lozano, A. C., Li, H., Niculescu-Mizil, A., Liu, Y., Perlich, C., Hosking, J., and Abe, N.: Spatial-temporal causal modeling for climate change attribution, in: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 587–596, 2009.

Luojus, K., Pulliainen, J., Takala, M., Derksen, C., Rott, H., Nagler, T., Solberg, R., Wiesmann, A., Metsamaki, S., Malnes, E., et al.: Investigating the feasibility of the GlobSnow snow water equivalent data for climate research purposes, in: Geoscience and Remote Sensing Symposium (IGARSS), 2010 IEEE International, IEEE, 2010.

Maddala, G. S. and Lahiri, K.: Introduction to econometrics, vol. 2, Macmillan New York, 1992.

Marinazzo, D., Pellicoro, M., and Stramaglia, S.: Kernel method for nonlinear Granger causality, Phys. Rev. Lett., 100, 144103, doi:10.1103/PhysRevLett.100.144103, 2008.

Martens, B., Miralles, D. G., Lievens, H., van der Schalie, R., de Jeu, R. A. M., Férnandez-Prieto, D., Beck, H. E., Dorigo, W. A., and Verhoest, N. E. C.: GLEAM v3: satellite-based land evaporation and root-zone soil moisture, Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-162, in review, 2016.

McCabe, M. F., Rodell, M., Alsdorf, D. E., Miralles, D. G., Uijlenhoet, R., Wagner, W., Lucieer, A., Houborg, R., Verhoest, N. E. C., Franz, T. E., Shi, J., Gao, H., and Wood, E. F.: The Future of Earth Observation in Hydrology, Hydrol. Earth Syst. Sci. Discuss., doi:10.5194/hess-2017-54, in review, 2017.

McCracken, M. W.: Asymptotics for out of sample tests of Granger causality, J. Econometrics, 140, 719–752, 2007.

McPherson, R. A., Fiebrich, C. A., Crawford, K. C., Elliott, R. L., Kilby, J. R., Grimsley, D. L., Martinez, J. E., Basara, J. B., Illston, B. G., Morris, D. A., Kloesel, K. A., Stadler, S. J., Melvin, A. D., Sutherland, A. J., Shrivastava, H., Carlson, J. D., Wolfinbarger, J. M., Bostic, J. P., and Demko, D. B.: Statewide monitoring of the mesoscale environment: A technical update on the Oklahoma Mesonet, J. Atmos. Ocean. Tech., 24, 301–321, doi:10.1175/JTECH1976.1, 2007.

Michaelsen, J.: Cross-validation in statistical climate forecast models, J. Clim. Appl. Meteorol., 26, 1589–1600, 1987.

Miralles, D. G., Holmes, T. R. H., De Jeu, R. A. M., Gash, J. H., Meesters, A. G. C. A., and Dolman, A. J.: Global land-surface evaporation estimated from satellite-based observations, Hydrol. Earth Syst. Sci., 15, 453–469, doi:10.5194/hess-15-453-2011, 2011.

Mokhov, I. I., Smirnov, D. A., Nakonechny, P. I., Kozlenko, S. S., Seleznev, E. P., and Kurths, J.: Alternating mutual influence of El-Niño/Southern Oscillation and Indian monsoon, Geophys. Res. Lett., 38, doi:10.1029/2010GL045932, 2011.

Mosedale, T. J., Stephenson, D. B., Collins, M., and Mills, T. C.: Granger causality of coupled climate processes: Ocean feedback on the North Atlantic Oscillation, J. Climate, 19, 1182–1194, 2006.

Myneni, R. B., Keeling, C. D., Tucker, C. J., Asrar, G., and Nemani, R. R.: Increased plant growth in the northern high latitudes from 1981 to 1991, Nature, 386, 698–702, 1997.

Nemani, R. R., Keeling, C. D., Hashimoto, H., Jolly, W. M., Piper, S. C., Tucker, C. J., Myneni, R. B., and Running, S. W.: Climate-Driven Increases in Global Terrestrial Net Primary Production from 1982 to 1999, Science (New York, N.Y.), 300, 1560–1563, doi:10.1126/science.1082750, 2003.

Nicholls, N. and Alexander, L.: Has the climate become more variable or extreme? Progress 1992–2006, Prog. Phys. Geog., 31, 77–87, 2007.

Papagiannopoulou, C., Miralles, D. G., Dorigo, W. A., Verhoest, N. E. C., Depoorter, M., and Waegeman, W.: Vegetation anomalies caused by antecedent precipitation in most of the world, Environ. Res. Lett., doi:10.1088/1748-9326/aa7145, 2017.

Pasini, A., Triacca, U., and Attanasio, A.: Evidence of recent causal decoupling between solar radiation and global temperature, Environ. Res. Lett., 7, 034020, doi:1748-9326-7-3-034020, 2012.

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

Rammig, A., Wiedermann, M., Donges, J. F., Babst, F., von Bloh, W., Frank, D., Thonicke, K., and Mahecha, M. D.: Coincidences of climate extremes and anomalous vegetation responses: comparing tree ring patterns to simulated productivity, Biogeosciences, 12, 373–385, doi:10.5194/bg-12-373-2015, 2015.

Rodriguez-Galiano, V. F., Ghimire, B., Rogan, J., Chica-Olmo, M., and Rigol-Sanchez, J. P.: An assessment of the effectiveness of a random forest classifier for land-cover classification, ISPRS Journal of Photogramm., 67, 93–104, 2012.

Rossow, W. and Duenas, E.: The international satellite cloud climatology project (ISCCP) web site: An online resource for research, B. Am. Meteorol. Soc., 85, 167–172, 2004.

Schneider U., Becker A. Finger P., Meyer-Christoffer A., Rudolf B., Ziese M. 2011 GPCC full data reanalysis version 6.0 at 0.5°: monthly land-surface precipitation from rain-gauges built on GTS-based and historic data, doi:10.5676/DWD_GPCC/FD_M_V7_050, 2011.

Seddon, A. W., Macias-Fauria, M., Long, P. R., Benz, D., and Willis, K. J.: Sensitivity of global terrestrial ecosystems to climate variability, Nature, 531, 229–232, 2016.

Shahin, M. A., Ali, M. A., and Ali, A. S.: Vector Autoregression (VAR) Modeling and Forecasting of Temperature, Humidity, and Cloud Coverage, in: Computational Intelligence Techniques in Earth and Environmental Sciences, Springer, 29–51, 2014.

Smith, T. M., Reynolds, R. W., Peterson, T. C., and Lawrimore, J.: Improvements to NOAA's historical merged land-ocean surface temperature analysis (1880–2006), J. Climate, 21, 2283–2296, 2008.

Stackhouse, Jr. Paul, W., Gupta, Shashi, K., Cox, Stephen, J., Mikovitz, C., Zhang, T., and Chiacchio, M.: 12-Year Surface Radiation Budget Dataset, GEWEX News, 14, 10–12, 2004.

Stock, J. H. and Watson, M. W.: Vector autoregressions, J. Econ. Perspect., 15, 101–115, 2001.

Su, L., Jia, W., Hou, C., and Lei, Y.: Microbial biosensors: a review, Biosens. Bioelectron., 26, 1788–1799, 2011.

Triacca, U.: Is Granger causality analysis appropriate to investigate the relationship between atmospheric concentration of carbon dioxide and global surface air temperature?, Theor. Appl. Climatol., 81, 133–135, 2005.

Tucker, C. J., Pinzon, J. E., Brown, M. E., Slayback, D. A., Pak, E. W., Mahoney, R., Vermote, E. F., and El Saleous, N.: An extended AVHRR 8-km NDVI dataset compatible with MODIS and SPOT vegetation NDVI data, Int. J. Remote Sens., 26, 4485–4498, 2005.

Tuttle, S. and Salvucci, G.: Empirical evidence of contrasting soil moisture–precipitation feedbacks across the United States, Science, 352, 825–828, 2016.

van der Werf, G. R., Randerson, J. T., Giglio, L., Collatz, G. J., Mu, M., Kasibhatla, P. S., Morton, D. C., DeFries, R. S., Jin, Y., and van Leeuwen, T. T.: Global fire emissions and the contribution of deforestation, savanna, forest, agricultural, and peat fires (1997–2009), Atmos. Chem. Phys., 10, 11707–11735, doi:10.5194/acp-10-11707-2010, 2010.

Verbesselt, J., Hyndman, R., Newnham, G., and Culvenor, D.: Detecting trend and seasonal changes in satellite image time series, Remote Sens. Environ., 114, 106–115, 2010.

Verbesselt, J., Umlauf, N., Hirota, M., Holmgren, M., Van Nes, E. H., Herold, M., Zeileis, A., and Scheffer, M.: Remotely sensed resilience of tropical forests, Nature Climate Change, 6, doi:10.1038/nclimate3108, 2016.

Von Storch, H. and Zwiers, F. W.: Statistical analysis in climate research, Cambridge university press, 2001.

Willmott, C. J. and Matsuura, K.: Terrestrial air temperature and precipitation: Monthly and annual time series (1950–1999), Center for climate research version, 1, 2001.

Wu, D., Zhao, X., Liang, S., Zhou, T., Huang, K., Tang, B., and Zhao, W.: Time-lag effects of global vegetation responses to climate change, Glob. Change Biol., 21, doi:10.1111/gcb.12945, 2015.

Xie, P. and Arkin, P. A.: Global precipitation: A 17-year monthly analysis based on gauge observations, satellite estimates, and numerical model outputs, B. Am. Meteorol. Soc., 78, 2539–2558, 1997.

Xie, P., Chen, M., Yang, S., Yatagai, A., Hayasaka, T., Fukushima, Y., and Liu, C.: A gauge-based analysis of daily precipitation over East Asia, J. Hydrometeorol., 8, 607–626, 2007.

Ye, L. and Keogh, E.: Time series shapelets: a new primitive for data mining, in: Proc. of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining – KDD '09, ACM Press, New York, USA, 947 pp., doi:10.1145/1557019.1557122, 2009.

Zeng, N., Hales, K., and Neelin, J. D.: Nonlinear dynamics in a coupled vegetation-atmosphere system and implications for desert-forest gradient, J. Climate, 15, 3474–3487, 2002.

Zhang, X., Alexander, L., Hegerl, G. C., Jones, P., Tank, A. K., Peterson, T. C., Trewin, B., and Zwiers, F. W.: Indices for monitoring changes in extremes based on daily temperature and precipitation data, Wiley Interdisciplinary Reviews: Climate Change, 2, 851–870, 2011.

Zwiers, F. W., Alexander, L. V., Hegerl, G. C., Knutson, T. R., Kossin, J. P., Naveau, P., Nicholls, N., Schär, C., Seneviratne, S. I., and Zhang, X.: Climate extremes: challenges in estimating and understanding recent changes in the frequency and intensity of extreme climate and weather events, in: Climate Science for Serving Society, Springer, 339–389, 2013.