



Inferring interdependencies from short time series

BEDARTHA GOSWAMI^{1,2,†}, PAUL SCHULTZ^{1,3,†,*}, BIRTE HEINZE¹, NORBERT MARWAN¹, BENJAMIN BODIRSKY¹, HERMANN LOTZE-CAMPEN^{1,4} and JÜRGEN KURTHS^{1,3}

¹Potsdam Institute for Climate Impact Research, P.O. Box 60 12 03, 14412 Potsdam, Germany

²Institute of Earth and Environmental Science, University of Potsdam, 14476 Potsdam, Germany

³Department of Physics, Humboldt University of Berlin, 12489 Berlin, Germany

⁴Department of Agricultural Economics, Humboldt University Berlin, 10099 Berlin, Germany

*Corresponding author. E-mail: pschultz@pik-potsdam.de

†These authors contributed equally to the research presented.

Abstract. Complex networks provide an invaluable framework for the study of interlinked dynamical systems. In many cases, such networks are constructed from observed time series by first estimating the interdependencies between pairs of datasets. However, most of the classic and state-of-the-art interdependence estimation techniques require sufficiently long time series for their successful application. In this study, we present a modification of the inner composition alignment approach (IOTA), correspondingly termed mIOTA, and review its advantages. Using two coupled auto-regressive stochastic processes, we demonstrate the discriminating power of mIOTA and show that it outperforms standard interdependence measures. We then use mIOTA to derive econo-climatic networks of interdependencies between economic indicators and climatic variability for Sub-Saharan Africa (AFR) and South Asia including India (SAS). Our analysis uncovers that crop production in AFR is strongly interdependent with the regional rainfall. While the gross domestic product (GDP) as an economic indicator in AFR is independent of climatic factors, we find that precipitation in the SAS influences the regional GDP, likely reflecting the influence of the summer monsoons. The differences in the interdependence structures between AFR and SAS reflect an underlying structural difference in their overall economies, as well as their agricultural sectors.

Keywords. Interdependence; correlation; inner composition alignment; time series analysis.

PACS Nos 05.45.; 02.50.-r; 02.50.Tt; 05.45.Tp; 02.70.Rr; 93.30.Bz; 93.30.Db; 92.60.Ry; 92.70.Kb; 92.70.Mn; 89.65.Gh; 89.75.Fb

1. Introduction

The theory of complex networks is a powerful tool that helps us to understand the interplay between local interactions and global features of dynamical systems [1]. The structure of the network which encodes all physically meaningful and relevant interactions is however, not clearly defined in all cases. In many situations, such as when studying interactions in the brain [2], climate [3], economic indicators [4], trade networks [5] or ecological populations [6], we are required to infer the structure of the network from a set of observed time series that encode the time evolution of the various system components. In such cases, the structure of the network, viz. the edges between the network nodes (represented here by the dynamical systems corresponding to each of the observed time series) are identified by first estimating all pairwise interdependencies and then ruling out those estimates that could

have occurred by random chance. The remaining set of interdependence estimates represents the edges of a complex network that encodes the non-random (assumably physically meaningful) interactions of the system under study.

Given the wide range of complexity of real-world systems, and the increasing number of datasets, several new approaches have been proposed over the last two decades that seek to extend our scope of inferring meaningful interdependencies from observed time series. The interdependence estimation approaches themselves stem from a broad spectrum of theoretical concepts such as information theory [7–9], symbolic dynamics [10, 11], recurrence plots [12–14], and Granger causality [15–17]. Notwithstanding the availability of such theoretically rigorous approaches, a common hindrance in the application of these techniques is that the interdependence estimation procedure is often unreliable for ‘short’ time series. This is particularly

relevant because, despite the increasing availability of data and the growing need of big data approaches, several real-world applications face limited data availability, such as very short time series or only few observations. Typical examples include paleoclimate proxy data [12], socio-economic measurements [18], gene regulatory networks in systems biology [19] or life threatening events such as epileptic seizures or ventricular tachycardia [20,21]. To the best of our knowledge, only a few approaches can cope with this challenge, such as recurrence-plot-based methods, and more recently, approaches such as that of Ernst *et al.* [22], Hempel *et al.* [23] and Ma *et al.* [24] that have been specifically designed to infer the interdependence between a pair of short time series.

The focus of this study is to provide an extension of the inner composition alignment (or IOTA) approach put forth by Hempel *et al.* [23] to infer interdependencies for short time series. In the following sections, we first review the ideas behind the IOTA approach and present its limitations, which motivate the extension presented in this study. Next, we outline the new approach, henceforth called modified inner composition alignment (mIOTA), and present its performance with respect to time series length, coupling strength, and noise. We also review its performance as a discriminating detector of interdependence for short time series in comparison to existing methods. Lastly, we apply our method to global socio-economic and climatic data sets, and ask the question: *What is the structure of interdependencies between economic indicators and climatic variability?* The purpose of the example is not to construct a *complex network* (as, for instance, in [25]) but to demonstrate how interdependencies can, in principle, be estimated from short time series to construct networks, opening up future research possibilities dealing with complex networks derived from short datasets.

We emphasise that the notion of what is a sufficiently ‘long’ (or ‘short’) time series depends crucially on the variability of the time series themselves: the higher the variability or stochasticity, the longer the time series required for a robust estimation of interdependence. That being said, for the rest of this study, we refer to time series with length ≤ 50 time points as ‘short’ time series.

2. Modified inner composition alignment (mIOTA)

2.1 Background and definition

Hempel *et al.* [23, 26] put forward an approach based on inner composition alignment (IOTA in short, denoted by the Greek letter ι), in order to infer the interdependence between a pair of short time series. According to

the authors, IOTA could be applied to very short time series up to only ten data points. Moreover, since the measure is asymmetric, it can be further used to infer the direction of the estimated interdependence as well.

The fundamental idea underlying IOTA is: Given there are two observables X and Y . In order to determine if X drives Y , we first estimate the permutation of X which rearranges the values of X in increasing order, and then apply the same permutation to the measurements of Y . If X drives Y , the values of Y will also be rearranged in an increasing order (see figure 1). If Y is independent of X , then the values of Y will be reordered in a random manner. The extent to which Y is reordered in monotonic increasing order is determined by the number of crossings that the reordered Y series has with itself. The lower the number of self-crossings, the more monotonic it is, and consequently, the more likely it is that it is driven by X .

Mathematically, let x_t and y_t , $t = 1, 2, 3, \dots, N$ be the observed time series of X and Y respectively (figure 1a), and $\pi_X(\cdot)$ be the permutation that arranges the entries x_t in increasing order. Let $y_t^{\pi_X} = \pi_X(y_t)$ be the permuted version of the time series y_t (figure 1b). Then the measure IOTA, ι , is given by

$$\iota^{(X \rightarrow Y)} = 1 - \frac{\sum_{i=1}^{N-2} \sum_{j=i+1}^{N-1} \omega_{ij} \Theta[(y_{j+1}^{\pi_X} - y_i^{\pi_X})(y_i^{\pi_X} - y_j^{\pi_X})]}{M}, \quad (1)$$

where ω_{ij} is an appropriate weight for each possible self-crossing, $\Theta(\cdot)$ denotes the Heaviside step function such that it is 1 for all positive inputs and zero otherwise, and $M = N(N-1)/2$ is the total number of possible self-crossings used here to normalise the observed number of self-crossings so that ι lies in the interval $[0, 1]$.

Even though, in principle, the normalisation M is applied to ensure that ι is in $[0, 1]$, in practice, ι values towards 0 are far rarer than expected. This is because of all the possible $N(N-1)/2$ self-crossings, some are more unlikely than others. Theoretically, it is possible to partially overcome this problem by an appropriate choice of weighting ω_{ij} , but in practice, it is hard to gauge how the weighting impacts the quality of estimation of interdependencies. This problem is highlighted with the surrogate distribution required to determine whether an estimated ι value has occurred purely by random chance (figure 1c). In most cases, we find that the distribution of ι values of completely uncorrelated pairs of time series are not centred around 0.5 as expected, but are rather biased towards 1, indicating that the value M is larger than is necessary. A serious consequence of this is that the approach stands to lose its discerning power between random and meaningful ι values as the surrogate distribution tends towards 1 and the interval of possibly

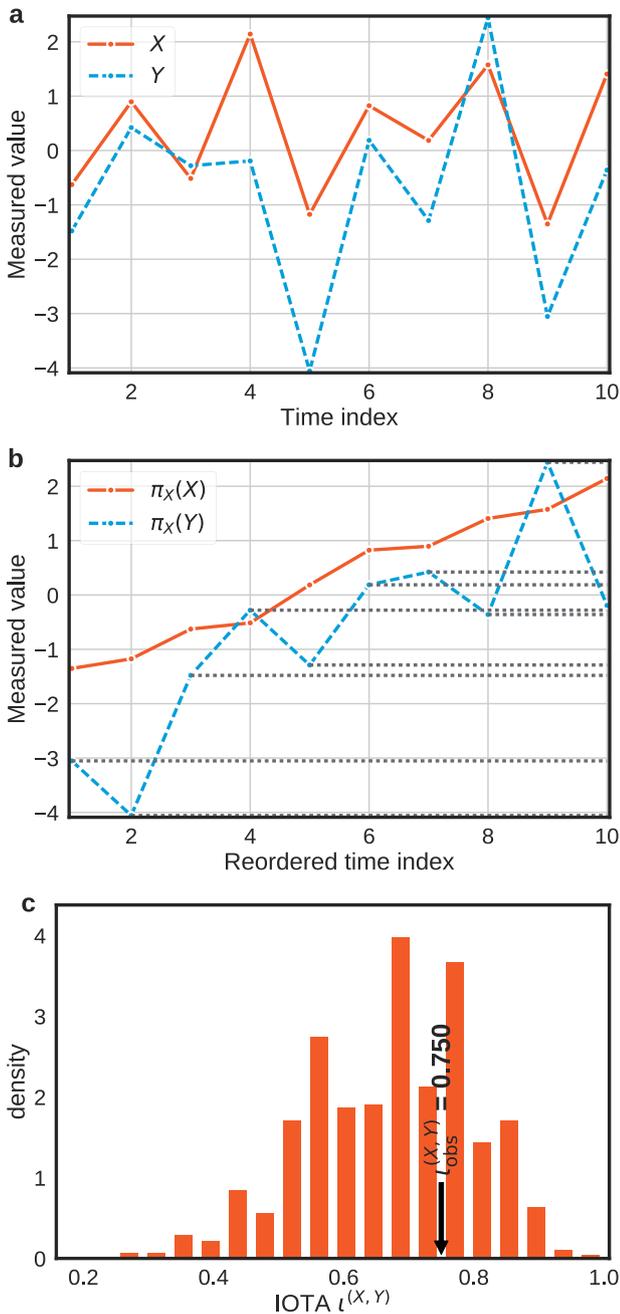


Figure 1. Inferring interdependencies with IOTA, ι . **(a)** Two sample Gaussian white noise time series X and Y of length 10 time points such that X drives Y . **(b)** The permuted versions $\pi_X(Y)$ (of Y) and $\pi_X(X)$ (of X) highlight that $\pi_X(Y)$ results in an almost monotonic reordering of Y as it is strongly driven by X . The reordering is not perfect as the two time series are not perfectly coupled. Horizontal dashed lines in **(b)** highlight the self-crossings of the reordered series $y_i^{\pi_X}$. **(c)** The surrogate distribution of ι values obtained by randomly shuffling y_i 10000 times and calculating $\iota^{X \rightarrow Y}$ for each randomisation. The observed $\iota_{obs}^{X \rightarrow Y}$ is indicated. The random surrogate distribution is not centred around 0.5, as would be expected from random time series, but is rather shifted towards the value of 1.

meaningful ι values gets smaller. Hempel *et al.* [26] recommends that the selection criterion for the detection of a statistically significant interdependence is that it should be greater than the expected value of the surrogate distribution. This however, only ensures that the observed ι is greater than what is expected by random chance – a much weaker null hypothesis than when trying to ensure that the observed value of a test statistic is less than 5% likely to have happened by random chance.

An additional drawback of the IOTA approach is that it fails to differentiate between positively correlated and negatively correlated interdependencies as all ι values are, by definition positive, and it is not clear what is the ι value of a pair of uncorrelated time series which would have ideally been the midpoint between perfectly correlated and perfectly anticorrelated ι values. In many situations, it is important to know not only that X is interrelated to Y , but also whether an increase in X leads to an increase in Y or not.

We propose to modify IOTA by counting the magnitude of increasing/decreasing pairwise distances in the permuted series $y_i^{\pi_X}$ instead of the number of self-crossings. Using the same notation as in eq. (1), we denote the modified IOTA (mIOTA) value μ as

$$\mu^{(X,Y)} = \frac{\sum_{i=1}^{N-1} \sum_{j=i+1}^N (y_j^{\pi_X} - y_i^{\pi_X}) \cdot |(y_j^{\pi_X} - y_i^{\pi_X})|}{\sum_{i=1}^{N-1} \sum_{j=i+1}^N (y_j^{\pi_X} - y_i^{\pi_X})^2}, \quad (2)$$

where $|\cdot|$ denotes the absolute value. The underlying idea here is that the higher the interdependence between X and Y , the more monotonic is $y_i^{\pi_X}$, and hence, the higher the magnitude of positive pairwise distances as we proceed along the reordered series $y_i^{\pi_X}$. The measured μ lies in the interval $[-1, 1]$. For a pair of uncorrelated time series, the sign of the increments $(y_j^{\pi_X} - y_i^{\pi_X})$ in eq. (2) will most likely be at random and thus the summands would approximately cancel each other out, resulting in a value close to zero. This is in contrast to adding up as would be assumed for correlated time series where $y_i^{\pi_X}$ is more monotonic. A positive value close to 1 indicates strong positive correlation, while a negative value close to -1 points to negative correlation.

Figure 2 schematically illustrates how μ helps to differentiate between positively and negatively correlated interdependencies for short time series. The random surrogate distribution for $\mu^{X \cdot}$ is centred around 0, unlike that of $\iota^{X \cdot}$, which helps to determine how unlikely it is to observe an extreme (positive or negative) value of μ by pure random chance. This advantage however comes with the trade-off that, unlike ι , the newly proposed measure μ is symmetric with respect to an interchange of X

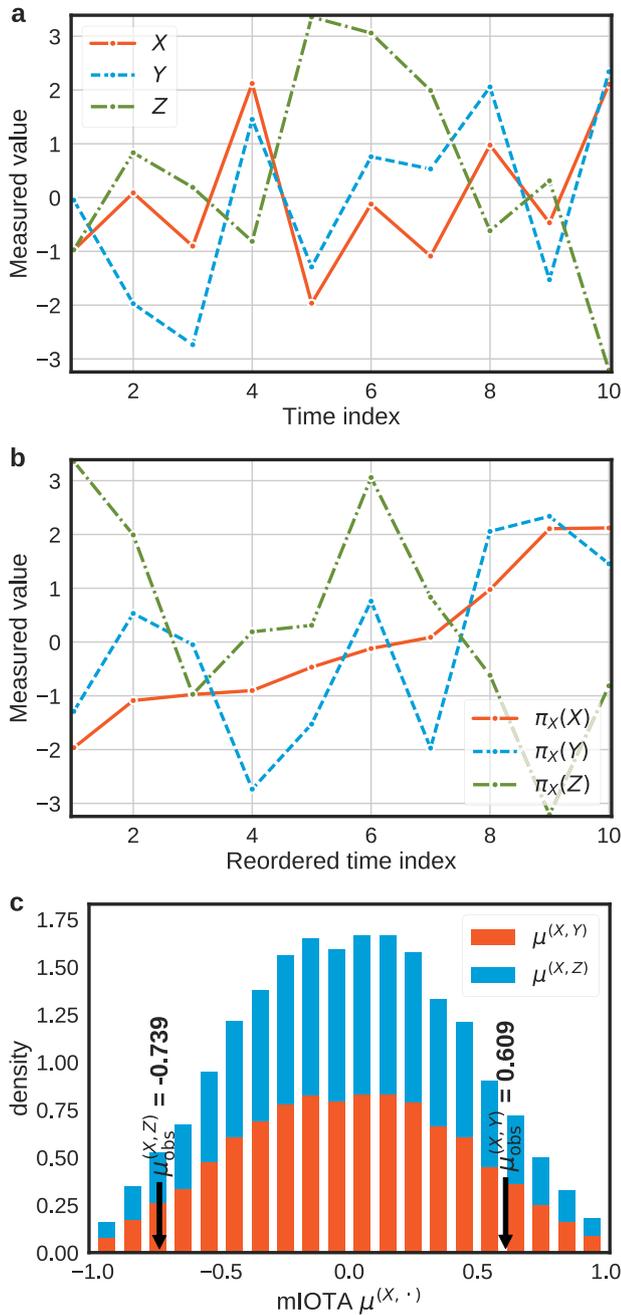


Figure 2. Inferring interdependencies with mIOTA, μ . (a). Three sample Gaussian white noise time series X , Y and Z are taken such that X drives Y in a positively correlated manner and X drives Z in a negatively correlated manner. (b) The permuted versions $\pi_X(Y)$ (of Y), $\pi_X(Z)$ (of Z), and $\pi_X(X)$ (of X). π_X reorders Y (Z) in increasing (decreasing) order, reflecting the underlying positive (negative) correlations between X and Y (X and Z). (c) Surrogate distributions of $\mu^{X,Y}$ ($\mu^{X,Z}$) values obtained by randomly shuffling y_t (z_t) 10,000 times and calculating $\mu^{X,Y}$ ($\mu^{X,Z}$) for each randomisation. The observed $\mu_{\text{obs}}^{X,Y}$ ($\mu_{\text{obs}}^{X,Z}$) is indicated. The random surrogate distributions are centred around 0, unlike in figure 1c, and the positive (negative) correlations between X and Y (X and Z) are well-differentiated from one another.

and Y (figure not shown) and thus cannot be used to infer the direction of the interdependence. The direction of the interdependence is different from the sign (positive or negative) of the correlation. The direction of interdependence is determined by whether X drives Y or vice versa, whereas the sign of correlation indicates whether an increase in the driver leads to an increase or decrease in the response. Although ι can in principle indicate whether X drives Y or not, it cannot tell us anything about the sign of the correlation. However, μ can differentiate whether X and Y are positively or negatively correlated, but it cannot tell us which one is the driven system and which one is the driver. Nevertheless, μ still offers a robust detection of interdependencies for short time series and performs better than existing methods. The details are discussed in the following subsections.

2.2 Predicting interdependencies between short, coupled AR(1) processes

To test the robustness of the proposed measure, we consider a system of coupled first-order auto-regressive processes without drift, henceforth referred to as AR(1) processes. The AR(1) processes chosen here is intended as a generic example to illustrate the performance of the proposed measure μ , as well as to help compare its performance with existing standard measures of quantifying interdependence. The proposed variant of the IOTA approach is at its core a statistical method to quantify predominantly linear interdependencies and hence can be applied to many model systems as well as to real-world datasets. The final conclusions are typically made based on a statistical significance test (outlined below) which helps identify those pairs of time series which show interdependence estimates outside reasonably chosen bounds corresponding to random chance.

The two AR(1) processes X and Y are generated such that X drives Y with a chosen coupling strength C ,

$$x_t = \frac{3}{4}x_{t-1} + \sigma\epsilon_t^X \quad (3)$$

$$y_t = \frac{3}{4}y_{t-1} + \sigma\epsilon_t^Y + Cx_t, \quad (4)$$

where $\epsilon_t^{(\cdot)}$ is a standard normal random variable with mean zero and variance one, and σ represents the magnitude of noise in the modelled AR(1) processes. For the sake of computational simplicity, we set the noise magnitudes and the AR(1) coefficient to the same values for both X and Y , with the understanding that the qualitative nature of our final conclusions is not impacted severely

by choosing different values for these parameters. The AR(1) time series described here are different from the Gaussian white noise time series used to schematically demonstrate the ideas underlying ι and μ in section 2.1.

The mIOTA approach is tested over changes in the time series length, coupling strength, and noise magnitude as follows:

1. For a given combination of time series length, coupling strength, and noise magnitude, we generate 100,000 realisations of X and Y (such that Y is driven by X).
2. For each of the above 100,000 (X, Y) combinations, we calculate the sample mIOTA, denoted here as $\mu_{\text{obs}}^{X,Y}$.
3. A surrogate distribution corresponding to the null hypothesis that the observed mIOTA values could have occurred by chance is constructed from 1000 random permutations of time series selected at random from the large set of Y time series. This is done by a randomisation of the timings of the observed Y time series. Thus, each surrogate time series is a random permutation of the corresponding observed time series of Y .
4. A two-sided p -value is obtained for each of the above 100,000 $\mu_{\text{obs}}^{X,Y}$ values based on the surrogate distribution.
5. Using a range of values of the false positive rate (FPR, which equals the significance level of the hypothesis test), we obtain a prediction result for each of the 100,000 values of $\mu_{\text{obs}}^{X,Y}$, which allows us to calculate the true positive rate (TPR) corresponding to each FPR.
6. The relation between TPR and FPR, also known as the receiver operating characteristic (ROC) curve is then used to estimate the area under the curve (AUC). The AUC attempts to quantify the predictive powers of the statistical test which tells us whether or not X drives Y . Ideally, we would like to have TPR values close to 1, irrespective of the value of FPR, which would lead to $\text{AUC} \approx 1$. A completely random prediction of whether X drives Y or not would, in theory, result in a TPR value approximately equal to the chosen FPR value, yielding an $\text{AUC} \approx 0.5$. We refer the readers to the discussions on ROC curves presented in Refs [27, 28].
7. We thus estimate an AUC value for each given combination of time series length, coupling strength, and noise magnitude.

AUC values near to or below 0.5 indicate bad discriminating power of the test, whereas AUC values close to one indicate a high discriminating power of the test. Our

results indicate that the mIOTA approach fares quite well for time series length of over 20 time points (figure 3). The efficacy of the measure is impacted by all three factors, but the length of the time series has a prominent effect.

2.3 Comparison of mIOTA (μ) with other measures

We repeat the AUC estimation as outlined in section 2.2 for the original IOTA ι , the proposed mIOTA μ , the classic Pearson’s cross-correlation coefficient [29,30] (Pearson’s ρ), Spearman’s rank-correlation coefficient [31] (Spearman’s ρ), and Kendall’s τ coefficient [32]. For all

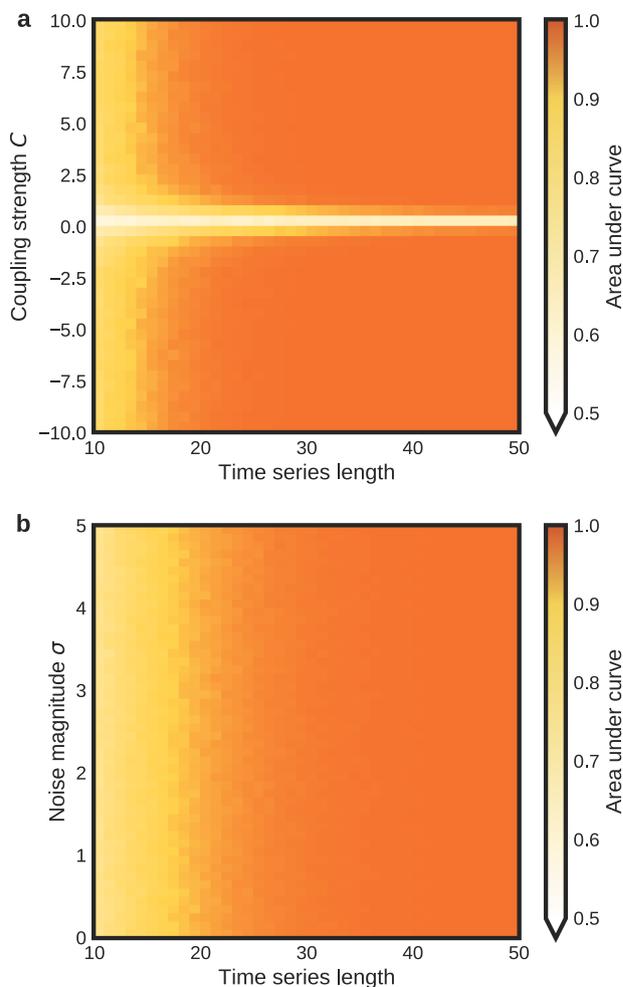


Figure 3. Predicting interdependencies for short, coupled AR(1) processes with mIOTA, μ . Using an ensemble of 10,000 coupled AR(1) processes X and Y according to eq. (3), we calculate $\mu_{\text{obs}}^{X,Y}$. The predictive performance is quantified via the AUC of a receiver operating characteristic, which we determined using time series surrogate distributions. This procedure is repeated for all combinations of the time series length and (a) coupling strength C or (b) noise magnitude σ . Details on the numerical approach are given in section 2.2.

the considered measures we estimate the p -value with the help of randomised surrogates (as detailed in the earlier section). The results (shown in figure 4) indicate that the mIOTA approach fares better in most of the combinations of time series length, coupling strength, and noise magnitude, where (as per figure 3) the AUC of μ is close to 1. We find that, for short time series, and for low coupling strengths, Pearson's ρ turns out to be the best measure for most combinations. This is in contrast to the widely held perception that Pearson's ρ is unsuitable for short time series. Kendall's τ is also seen to perform well in a significant number of combinations, although without any discernible relation to the

parameter combination – an interesting insight considering the widely held perception that it is most suited for inferring interdependencies in short time series. We note however, that even though we have considered a large ensemble size, the seemingly random scatter of the wins for Kendall's τ might be due to sample size effects. It also becomes clear from figure 4 that the earlier definition of IOTA does not fare well in most of the parameter combinations chosen in the analysis.

Note that in figure 4a, barring the random scatter for the wins by Kendall's τ , mIOTA is the best of the chosen set of measures for time series with length over 30 time points. Combining this insight with the AUC values for time series length ≥ 30 from figure 3, we conclude that mIOTA is a good choice for estimating interdependencies for short time series with 30–50 time points.

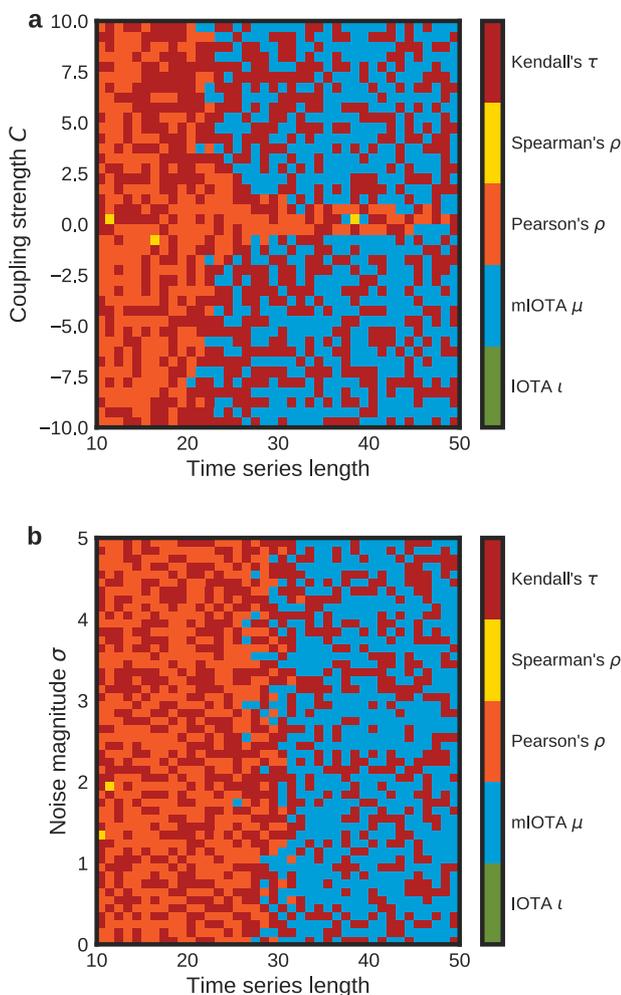


Figure 4. Comparison of mIOTA with other approaches. For comparing mIOTA μ with IOTA ι , Pearson's ρ , Spearman's ρ and Kendall's τ coefficient, we ranked them according to their predictive performance measured by the respective AUC. Here, we depict the correlation measure with highest AUC for all combinations of the time series' length and (a) coupling strength C or (b) noise magnitude σ . The numerical approach for each measure is analogous to the one yielding figure 3; details are given in section 2.3.

3. Interdependencies between economic indicators and climatic variability

3.1 Data

We use historic economic and climatic data that are input data for the framework of the global land use allocation model MAGPIE (Model of Agricultural Production and its Impact on the Environment) [33] which is linked to the land-use model LPJmL [34]. The data has been aggregated to 10 world regions which were clustered based on similar socio-economic, cultural and climatic conditions as well as geographical proximity and comparable size. We base our analysis on the economic world regions in MAGPIE as a compromise between the very detailed grid level of climate observables – sharing common properties within a climate zone – and socio-economic indicators, where information is aggregated only on a national level. The MAGPIE regions are characterised not only on the basis of economic indicators but also the climate. Hence, we assume that interdependencies in the data aggregated to a regional level are also valid for a majority of constituting subregions.

Here, we present the results of interdependence estimations using mIOTA μ on two of the ten MAGPIE regions: (i) Sub-Saharan African (AFR), and (ii) South Asia including India (SAS). The economic and demographic indicators that we extract from the MAGPIE database include: (i) crop production (tons), (ii) crop demand, (iii) livestock production, (iv) livestock demand, (v) population numbers (1000 heads), and (vi) gross domestic product (GDP) (US\$). These datasets are resolved annually for a period of 46 years (1961–2007), and are available on a per-country basis. For the purpose

of our analysis, we further aggregate the time series so that we have one time series for each of the above indicators for each of AFR and SAS. The aggregation is done by summing the total value for each of the above indicators within each respective MAGPIE region.

In addition to the economic indicators listed above, we consider climatic indicators: the annual mean temperature ($^{\circ}\text{C}$) and the annual sum of precipitation (mm/yr) on a grid level given ca. 60.000 cells ($0.5^{\circ}\text{lon} \times 0.5^{\circ}\text{lat}$). We extract the climatic variables for the AFR and SAS regions from the globally gridded dataset, and further aggregate them to obtain (as with the economic time series) one time series each for temperature and precipitation for each of AFR and SAS. The aggregation is done by taking the average annual temperature and the total precipitation sum for each of the two MAGPIE regions considered.

3.2 Estimation of μ

We first remove the influence of long-term trends by considering only the rate time series, i.e., the time series of *changes per year* in all of the datasets considered (both economic and climatic). Next, since the data include disparate variables with different units, we normalise them using a quantile-based normalisation so that the transformed time series approximately follow a standard normal distribution. That is, for a given time series x_t , we obtain the transformed time series x_t^* as

$$x_t^* = \Phi^{-1}\left(1 - \frac{q(x_t)}{100}\right), \quad (5)$$

where $\Phi^{-1}(\cdot)$ denotes the inverse of the cumulative distribution function of a standard normal distribution, and $q(\cdot)$ is the quantile function that returns the quantile of each given value of x_t . This transformation does not change the relative ordering of the measurements in a time series, which is crucial for the estimation of μ .

We thus have, for each of AFR and SAS, the following time series:

1. Normalised rates of total regional crop production ΔCP_t^* .
2. Normalised rates of total regional crop demand ΔCD_t^* .
3. Normalised rates of total regional livestock production ΔLP_t^* .
4. Normalised rates of total regional livestock demand ΔLD_t^* .
5. Normalised rates of total regional gross domestic product ΔGDP_t^* .

Similarly, for the climatic variables, we have for each of AFR and SAS, the following two datasets:

1. Normalised rates of average regional temperature ΔT_t^* .
2. Normalised rates of total regional precipitation ΔP_t^* .

These normalised datasets are visualised for AFR in figure 5 (upper panel) and for SAS in figure 6 (upper panel). For each of these two regions, we thus have 7 datasets leading to a possible number of 21 interdependence links between them. We estimate μ for each of these 21 combinations. Next, for each combination, we randomise the data 10,000 times, and estimate μ for each random pair, allowing us to construct a surrogate distribution corresponding to the null hypothesis: *The considered pair of time series is not interdependent*. The surrogate distribution is used to determine whether the observed value of μ between the actual time

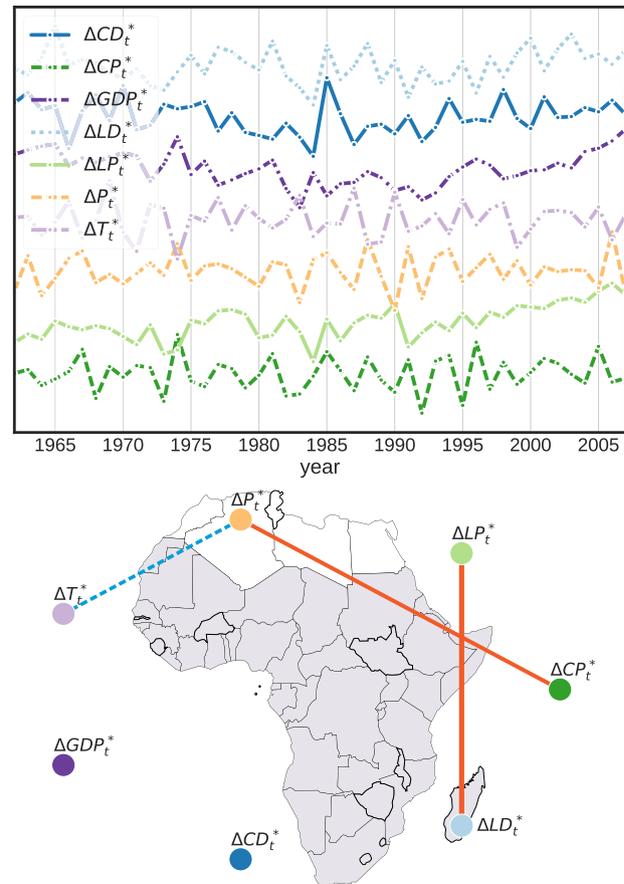


Figure 5. Climate-economic interdependencies in Sub-Saharan Africa (AFR). (**upper panel**) Schematic representation of the normalised rate time series. The time series are separated along the ordinate for visual guidance. (**lower panel**) Statistical interdependence network for AFR. The node colour and label correspond to the time series in the upper panel. The link width is proportional to the value of mOTA μ , while the link colour is solid orange for positive values of μ and dashed blue otherwise [35].

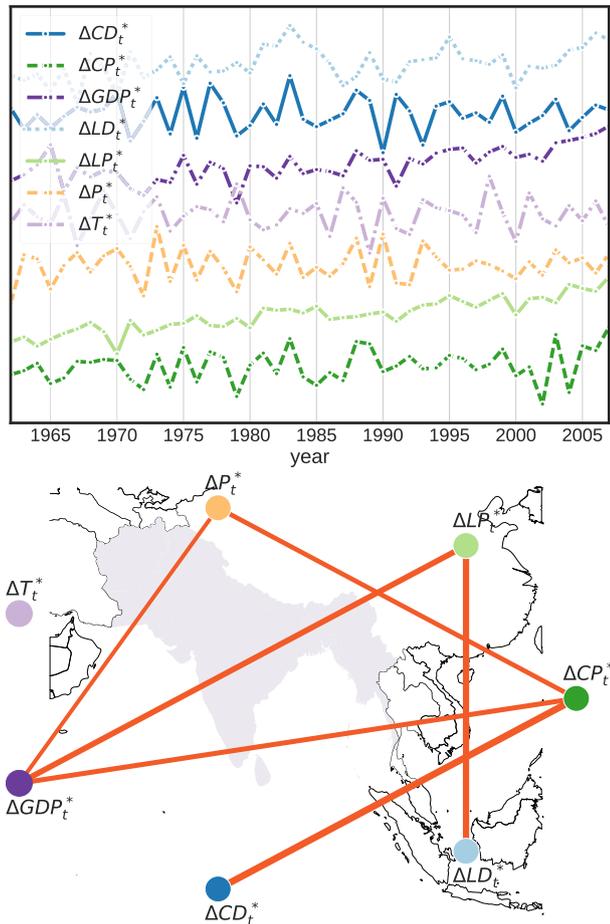


Figure 6. Climate-economic interdependencies in South Asia (SAS). (**upper panel**): Schematic representation of the normalised rate time series. The time series are separated along the ordinate for visual guidance. (**lower panel**) Statistical interdependence network for SAS. The node colour and label correspond to the time series in the upper panel. The link width is proportional to the value of mIOTA μ , while the link colour is solid orange for positive values of μ and dashed blue otherwise [36].

series (not randomised) is statistically significant with respect to the null hypothesis at a confidence level of 5% (which is also the FPR of the test). The combinations which result in significantly lower or higher values of μ (based on a two-tailed test) are visualised as a weighted interdependence edge between the corresponding economic/climatic observables in figure 5 (lower panel) for AFR and in figure 6 (lower panel) for SAS.

3.3 Econoclimatic networks

Our analysis reveals that crop production in the AFR region is significantly interlinked only to the total regional precipitation ($\mu_{\text{AFR}}^{\Delta P, \Delta CP} = 0.55$). In the con-

text of growing concern of projected crop yield changes under changing climatic conditions in Africa [37], this result indicates that out of regional temperature and regional precipitation, changes in the latter holds more influence over future crop yields. It is interesting to note further that the regional crop production could not be linked significantly to the regional crop demand, a possible reason for which could be that international trade balances out production-side fluctuations. Contrarily, we also find that the AFR livestock production is significantly interdependent to its demand ($\mu_{\text{AFR}}^{\Delta LD, \Delta LP} = 0.7$). We further find that the GDP in AFR is not significantly linked to either crop production or livestock production, indicating additional economic factors that contribute to the GDP.

The analysis further indicates a significant negative interdependence between the two climatic variables – regional mean temperature and the total regional precipitation – with $\mu_{\text{AFR}}^{\Delta T, \Delta P} = -0.49$.

In the SAS region, we find that, out of the two climatic variables, only precipitation is linked to the economic indicators. Precipitation in this region is positively correlated to crop production and GDP. This is likely due to the strong influence of the Asian Summer Monsoon (ASM) on the region’s agriculture and also on the fact that agricultural production forms a significant part of the economic turnover of the countries in this region. In contrast to the situation in AFR, the GDP in SAS is interdependent on crop production and livestock production, besides precipitation. This difference might be indicative of an underlying structural difference in the overall economies between the AFR and the SAS countries.

Regional crop production in SAS is further linked to its regional demand, which highlights a further dissimilarity with AFR, where crop production was seen to be decoupled from its demand. As with the overall economy, this would indicate that the agricultural economies in these two regions are also structurally different: SAS seems to be less connected to international agricultural markets than AFR. This indicates, however, that consumers in SAS might be more exposed to climatic variations as they make less use of the buffer capacity of international markets.

Interestingly, we find that neither livestock production nor its demand in the SAS is linked to either crop production or crop demand. This is contrary to the perception that livestock has a crucial impact on the agricultural economy of India (one of the major countries of the region). The absence of links between datasets, however, does not necessarily imply that *there are no interdependencies between the two*. Rather, our results reflect that *given the datasets, and the best performing interdependence measure μ , our analysis could not reject*

the null hypothesis that there was no interdependence between these combinations of time series. The reasons might be that the livestock sector is too small compared to plant-based food consumption to have a significant effect or that ruminant production is still mainly grass- and scavenging-based and therefore has little interaction with the cropping sector.

The link between precipitation and GDP in SAS may also be an example of a *spurious* link, likely caused here by an indirect correlation of precipitation on the GDP *via crop production*. This is not to suggest that precipitation *does not* influence GDP, but rather that precipitation might not influence GDP *directly*. The problem of spurious links is a general problem faced by all interdependence estimates, and is caused by either an *indirect influence* or a *common driver*. Measures such as partial cross-correlation, conditional mutual information, and more recently, an approach based on transfer entropy [38], and another that uses a recurrence-based measure of conditional dependence [14] seek to overcome this problem by ruling out indirect and common driver influences. The mIOTA approach itself needs to be further extended to be able to rule out spurious links in short time series.

In summary, our results highlight those pairs of econoclimatic variables which are quite likely to be interdependent on each other. The results aim to supplement, and they themselves need to be supplemented by, a thorough understanding of the underlying mechanisms which moderate these interrelations.

4. Conclusion

We have presented an extension of the IOTA approach to detect interdependencies for short time series. The proposed modification of IOTA, or mIOTA, overcomes the drawbacks that IOTA is unable to distinguish between positive and negative correlations, and that the null distribution for IOTA is biased towards higher values. The improvements in mIOTA come with the trade-off that, unlike IOTA, it is unable to uncover the direction of interdependencies. With a simple example of a pair of coupled AR(1) processes, by varying the coupling strength, time series length, and the magnitude of noise, we have demonstrated that mIOTA outperforms not only the earlier version of IOTA, but also standard interdependence measures such as Pearson's cross-correlation coefficient, Spearman's rank-correlation coefficient, and Kendall's tau. We further stress that, although the proposed mIOTA approach does not quantify causal relations (unlike IOTA, or Granger causality), based on the comparative results presented here, we find that mIOTA does well to indicate

interrelations, especially for short time series. This is at the cost of not being able to indicate the causal relation of an interdependence.

For sufficiently long time series, a sliding window analysis of mIOTA values can be carried out to estimate dynamically changing interactions between two observables. Calculating mIOTA for each window results in a time series of mIOTA values which encapsulates the time evolution of the interdependence between the two observables. Again, appropriate statistical tests can be designed to test whether the obtained mIOTA values could have been obtained by random chance. Notably, a sliding window analysis with mIOTA can be performed on smaller windows (i.e. shorter time series) than with competing measures, leading to a finer resolution of interdependence changes.

We used mIOTA to construct econoclimatic networks between climatic observables (mean regional temperature and total regional precipitation) and economic indicators (crop production, crop demand, livestock production, livestock demand, and GDP) for Sub-Saharan Africa and South Asia including India. Our analysis reveals that climate (*viz.* precipitation) plays an important role in crop and livestock production in the AFR region but does not influence its GDP. Precipitation in the SAS, however, is found to influence the regional GDP, most likely reflective of the influence of the Asian summer monsoon. Differences in the interdependence structures between the two regions hint at underlying structural difference in their economies.

Our study places the construction of networks from short time series on a robust footing, and opens up future research possibilities that might aim to derive complex networks from short datasets. We also note that mIOTA has the potential to be further improved so that it can possibly rule out spurious interrelations induced either by a common driver or by an indirect interdependence. This remains the subject of ongoing and future work for us.

Acknowledgements

This paper was developed within the scope of the IRTG 1740/TRP 2011/50151-0, funded by the DFG/FAPESP, and the DFG project IUcliD (DFG MA4759/8). BG was partially supported by the MWFK Brandenburg. PS gratefully acknowledges the support of BMBF, CoNDyNet, FK. 03SF0472A. NM and BG have received funding from the European Unions Horizon 2020 Research and Innovation programme under the Marie Skłodowska-Curie grant agreement No 691037 (project QUEST). We thank Christian Stegmaier and Markus Bonsch for their contributions in assembling and

validating the datasets. BG wishes to thank Aljoscha Rheinwalt for the helpful discussion on random sampling and sample size effects.

References

- [1] S Boccaletti, V Latora, Y Moreno, M Chavez and D-U Hwang, *Phys. Rep.* **424**, 175 (2006)
- [2] Ed Bullmore and O Sporns, *Nat. Rev. Neurosci.* **10**, 186 (2009)
- [3] J F Donges, Y Zou, N Marwan and J Kurths, *Eur. Phys. J. Spec. Top.* **174**, 157 (2009)
- [4] G Bonanno, G Caldarelli, F Lillo and R N Mantegna, *Phys. Rev. E* **68**, 046130 (2003)
- [5] J Maluck and R V Donner, *Eur. Phys. J. B* **90**, 26 (2017)
- [6] J A Steele, *et al. ISME J.* **5**, 1414 (2011)
- [7] T Schreiber, *Phys. Rev. Lett.* **85**, 461 (2000)
- [8] K Hlaváčková-Schindler, M Paluš, M Vejmelka and J Bhattacharya, *Phys. Rep.* **441**, 1 (2007)
- [9] B Pompe and J Runge, *Phys. Rev. E* **83**, 051122 (2011)
- [10] J I Deza, M Barreiro and C Masoller, *Eur. Phys. J. Spec. Top.* **222**, 511 (2013)
- [11] A Müller, J F Kraemer, T Penzel, H Bonnemeier, J Kurths and N Wessel, *Physiol. Meas.* **37**, R46 (2016)
- [12] N Marwan, M H Trauth, M Vuille and J Kurths, *Clim. Dyn.* **21**, 317 (2003)
- [13] J H Feldhoff, R V Donner, J F Donges, N Marwan and J Kurths, *Phys. Lett. A* **376**, 3504 (2012)
- [14] A M T Ramos, A Builes-Jaramillo, G Poveda, B Goswami, E E N Macau, J Kurths and N Marwan, *Phys. Rev. E* **95**, 052206 (2017)
- [15] Y Chen, G Rangarajan, J Feng and M Ding, *Phys. Lett. A* **324**, 26 (2004)
- [16] C Diks and M Wolski, *J. Appl. Econ.* (2015)
- [17] D A Smirnov, N Marwan, S F M Breitenbach, F Lechleitner and J Kurths, *Europhys. Lett.* **117**, 10004 (2017)
- [18] J Reyes, S Schiavo and G Fagiolo, *J. Int. Trade Econ. Dev.* **19**, 215 (2010)
- [19] H de Jong, *J. Comput. Biol.* **9**, 67 (2002)
- [20] K Lehnertz, R G Andrzejak, J Arnhold, T Kreuz, F Mormann, C Rieke, G Widman and C E Elger, *J. Clin. Neurophysiol.* **18**, 209 (2001)
- [21] A Schumann, N Wessel, A Schirdewan, K J Osterziel and A Voss, *Stat. Med.* **21**, 2225 (2002)
- [22] J Ernst, G J Nau and Z Bar-Joseph, *Bioinformatics* **21**(suppl 1), i159 (2005)
- [23] S Hempel, A Koseska, J Kurths and Z Nikoloski, *Phys. Rev. Lett.* **107**, 054101 (2011)
- [24] H Ma, K Aihara and L Chen, *Sci. Rep.* **4**, 7464 (2014)
- [25] K Rehfeld, N Marwan, S F M Breitenbach and J Kurths, *Climate dynamics* **41**, 3 (2013)
- [26] S Hempel, A Koseska and Z Nikoloski, *Eur. Phys. J. B.* **86**, 250 (2013)
- [27] J A Hanley and B J McNeil, *Radiology* **143**, 29 (1982)
- [28] S J Mason and N E Graham, *Q. J. R. Meteorol. Soc.* **128**, 2145 (2002)
- [29] K Pearson, *Proc. Roy. Soc. London* **58**, 240 (1895)
- [30] S M Stigler, *Stat. Sci.* **4**, 73 (1989)
- [31] C Spearman, *The American journal of psychology* **15**, 72 (1904)
- [32] M G Kendall, *Biometrika* **30**, 81 (1938)
- [33] H Lotze-Campen, C Müller, A Bondeau, A Jachner, A Popp and W Lucht, *Agric. Econ.* **39**, 325 (2008)
- [34] A Bondeau, *et al. Global Change Biol.* **13**, 679 (2007)
- [35] The map of Sub-Saharan Africa has been modified from http://commons.wikimedia.org/wiki/File:Africa_map_regions.svg, which is under public domain.
- [36] The map of South Asia has been modified from Tauhid16, https://commons.wikimedia.org/wiki/File:Members_of_BIMSTEC.svg, which is under the attribution-share-alike 4.0 international license. the license terms can be found on the following link: <https://creativecommons.org/licenses/by-sa/4.0/>.
- [37] J Knox, T Hess, A Daccache and T Wheeler, *Environ. Res. Lett.* **7**, 034032 (2012)
- [38] J Runge, J Heitzig, V Petoukhov and J Kurths, *Phys. Rev. Lett.* **108**, 258701 (2012)