Spatio-temporal patterns of diurnal temperature: a random matrix approach I-case of India

Madhuchhanda Bhattacharjee Department of Mathematics University of Manchester chhanda.bhatta@gmail.com Arup Bose^{*} Statistics and Mathematics Unit Indian Statistical Institute bosearu@gmail.com

Abstract

We consider the spatio-temporal gridded daily diurnal temperature range (DTR) data across India during the 72-year period 1951–2022. We augment this data with information on the El Niño-Southern Oscillation (ENSO) and on the climatic regions (Stamp's and Koeppen's classification) and four seasons of India.

We use various matrix theory approaches to trim out strong but routine signals, random matrix theory to remove noise, and novel empirical generalised singularvalue distributions to establish retention of essential signals in the trimmed data. We make use of the spatial Bergsma (S_B) statistics to measure spatial association and identify temporal change points in the spatial-association.

In particular, our investigation captures a yet unknown change-point over the 72 years under study with drastic changes in spatial-association of DTR in India. It also brings out changes in spatial association with regard to ENSO.

We conclude that while studying/modelling Indian DTR data, due consideration should be granted to the strong spatial association that is being persistently exhibited over decades, and provision should be kept for potential change points in the temporal behaviour, which in turn can bring moderate to dramatic changes in the spatial association pattern.

Some of our analysis also reaffirms the conclusions made by other authors, regarding spatial and temporal behavior of DTR, adding our own insights. We consider the data from the yearly, seasonal and climatic zones points of view, and discover several new and interesting statistical structures which should be of interest, especially to climatologists and statisticians. Our methods are not country specific and could be used profitably for DTR data from other geographical areas.

Keywords: Bergsma's correlation, climatic regions, correlation matrix, diurnal temperature range (DTR), empirical spectral distribution, empirical generalised singular-value

^{*}Research supported by J.C. Bose National Fellowship, JBR/2023/000023 from Science and Engineering Board, Govt. of India.

distribution, El Niño Southern oscillation, generalised singular-value decomposition, high dimensional time series, Marčenko-Pastur law, Pearson correlation, singular-value decomposition, spatial association measure, spatial dependence.

AMS Subject classification: Primary 62P12; Secondary 62H11, 62M10, 62H20, 62G05.

1 Introduction

Let T_{max} and T_{min} denote the maximum and minimum near surface air temperatures within a day (a period of 24 hours). Then the *Diurnal Temperature Range* (DTR) for that day is defined as $T_{max} - T_{min}$. Recorded DTR data over two centuries across the globe has become an increasingly important spatio-temporal historical record, especially due to global warming. It is acknowledged that changes in many environmental indices are causally linked to changes in local and global DTR values. Variations in the DTR has significant impact on life and agriculture. It is important to understand the space-time behaviour of DTR across locations over time.

Time series of *yearly average* DTR from different regions across the world have been investigated by several authors from various perspectives, Zhong et al. [2023], Sharma et al. [2021], Stjern [2020], Mall et al. [2021], Sen-Roy [2019], to mention a few. Some of these articles study DTR at global level (for example Zhong et al. [2023] and Stjern [2020]), while others make specific contributions on knowledge of DTR for the Indian subcontinent (for instance Sharma et al. [2021] Stjern [2020], Mall et al. [2021], Sen-Roy [2019], Vinnarasi et al. [2017], Jhajharia and Singh [2011] and Kothawale et al. [2010]).

All of these studies primarily focus on the *temporal pattern* of DTR. A few do recognise the existence of spatial variation (that is, variations across space) and present some results at individual spatial units, like grids. Terminologies such as "spatial distributions" and "spatial variations" that appear in some of these articles, always are synonymous with presentation of the results in the form of a map. It is typically left to the reader to gain any insight into possible "spatial" variation. Since the DTR is measured at specific times and locations, terms like "spatial distribution" might be a bit misleading, since that would normally refer to some randomness in the space variable. The only randomness here is in the magnitude of the DTR.

The randomness in DTR is expected to have strong temporal as well as spatial components. However the literature appears to have focused extensively on studying the temporal behaviour of DTR, without attempting to capture the inherent spatial dependence. Any such attempt should involve the spatial information of the grids.

Admittedly the temporal component has a strong presence which masks the other contributors like spatial information. Wu et al. [2009] developed *Multidimensional Ensemble Empirical Mode Decomposition* (MEEMD) and this was applied by Ji et al. [2014] to climate data for eliminating the oscillatory component of a time series and potentially reveal the slow varying components. Vinnarasi et al. [2017] used this technique to Indian DTR data. However, such techniques do not account for interdependence of the DTR series arising from various geographic locations.

We consider the DTR data for India, which is available since 1951. We try to extract the interdependence of the spatial units with regard to DTR. Since the common temporal factors make it difficult to isolate such salient patterns, we employ novel linear (random) matrix theory techniques to de-trend the data temporally, while preserving essential spatial signals. This technique allows us to study the time series jointly rather than detrend them individually. We employ a second novel technique to quantify the extent of spatial (statistical) dependence by a univariate statistics. This eliminates the need of having to rely on subjective judgements based on maps to arrive at an understanding of the degree of spatial association/dependence. Our method allows implementation at various temporal units and windows, which enables us to study the temporal change in such spatial dependence. This lead us to uncover yet unknown striking changes in spatial association pattern of DTR across India.

We emphasise that though we focus on India, our methods are not region specific and are applicable mutatis mutandis to other regions of the world. The analysis pipeline proposed here can be implemented using standard software like **R** for any similar cohort of time series that arise from a collection of geographical units.

Section 2 on materials and methods is divided into two sub-sections. In Section 2.1, after describing the data, which is spatio-temporal in nature, we discuss the issues related to organising it appropriately in a matrix form. Then in Section 2.2 we describe in brief the ideas that we have used from matrix and random matrix theory. These include the regular and generalised singular-value decompositions, empirical spectral distribution, Marčenko-Pastur law. We also describe Spatial-Bergsma statistic which we will use as a statistical measure of spatial association.

In Section 3 we present the result of our preliminary explorations using the above ideas. In Section 4 we dig into a very detailed analysis, and describe the results of our analysis. In Section 5 we summarise our findings briefly.

In Section S1 we show some additional analysis of the DTR, especially under different stratifications using climatic zones and seasons. Since there are several figures in this article, in Section S2 we provide at one place the descriptions of all the figures.

2 Materials and methods

2.1 Materials: the data

In Figure 1, the spatial domain of India is split into a grid of 362 locations, spanning six climatic regions, between $7.5^{\circ}N-37.5^{\circ}N$ and $67.5^{\circ}E-97.5^{\circ}E$, with a spatial resolution of $1^{\circ} \times 1^{\circ}$. The daily DTR data at each location for the 72-year period 1951–2022 is sourced from Climate Research Services, India Meteorological Department, Govt. of India (*https://www.imdpune.gov.in/lrfindex.php*). Thus, we have 362 daily time series spanning 72 years, with some missing values. Complete data is available only for 280 of the 362 locations. This data could be grouped according to months/seasons/years in temporal dimension and/or according to the six climatic regions in the spatial dimension. Such grouping will of course depend on the aspects that we wish to probe.



Figure 1: Climate grid of India: 1-Tropical monsoon; 2-Tropical savannah, wet and dry; 3-Arid, steppe, hot; 4-Humid subtropical; 5-Montane climate; 6-Hot deserts, Arid.

2.1.1 Data matrix

The time series of the daily DTR values for the 280 (or 362) locations for the 72 years yields a 280 (or 362) dimensional vector time series with 26298 time points. We visualise this as a 26298 × 280 (or 362) large *data matrix* X say. For a specific year, the row dimension for the data sub-matrix is 365 or 366, and the rows are arranged in the naturally increasing order of time. For now, the 280 (or 362) locations (columns) are arranged in some fixed order (e.g. latitude x longitude), and we shall have more to say on this below. The entries of X are subject to location dependent temporal patterns (e.g. seasonal variation). There is also spatial dependence across the locations, which is expected to be non-stationary across different pairs of locations.

2.1.2 Ordering the locations

Naively, the collection of 280 or 362 time series can be simply put side-by-side to form a data matrix for analysis. However each series arising out of a grid also has two key additional pieces of information, namely their latitudes and longitudes. Therefore, to investigate the temporal and/or spatial behaviour of the DTR across times and regions, we need to decide how this two-dimensional location information should be incorporated so as to arrive at an appropriate ordering amongst these time series. The consequent visual message might (and does) depend on this arrangement.

There are several choices for ordering the locations, dictated by convenience, geographical proximity, or division of India into the six climatic zones. One may begin with the lowest latitude, and move according to the increasing order of longitude of the locations, then come back to the next latitude and repeat the process. One could even switch the roles of the longitudes and latitudes. In any case, with this approach, two regions in the same climatic zone might end up being distant in the ordering.

On the other hand, we could use a classification of the locations into climatic zones. The Köppen [1918] classification of climatic zones of the world was later modified by several scientists (see McKnight and Hess [2000]). This yields a classification of India into *six* major climatic zones as given in Figure 1. We could group the locations according to the climatic zones they belong to. This would imply that locations that are geographically close, are not necessarily close with respect to the ordering. One may use the latitude-longitude ordering within each zone to alleviate this problem somewhat.

We tried several different orderings, and eventually settled on a *spiral ordering* that is a type of Hilbert space-filling curve, which are sometimes used in image processing to rearrange image pixels in order to enhance pixel locality. Mathematically this can be thought of as a bijection from \mathbb{N}^2 to \mathbb{N} , where $\mathbb{N} = \{1, 2, \ldots\}$. In our context this addresses both the zonal and spatial proximity of the grids reasonably well.

Visualising the locations according to the longitudes and latitudes, we start with the lowest latitude and longitude position, say (1,1) position, then move to the (1,2)position, and then move down diagonally to the (2,1) position, then move down to the (3,1) position, then successively to (2,2), (1,3), and so on. In other words, we move anti-diagonally and wrap around. Likewise, instead of first moving to (1,2), one could move first to (2,1) etc. We refer to this ordering as the *spiral ordering*. This ordering provides a significantly improved visual depiction and will get reflected in our figures and plots, thereby conveying valuable insight into the data and its behaviour from the spatial point of view. We shall call this 26298 × 280 reorganised data matrix D.

2.2 Methods

We have used ideas from linear matrix theory, random matrix theory (RMT), and statistics to analyse the data. These include the regular and generalised singular-value decompositions (SVD and GSVD), empirical spectral distribution (ESD), Marčenko-Pastur (MP) law, Bergsma's correlation ρ , and Spatial-Bergsma (S_B) which is a statistical measure of spatial association based on Bergsma's correlation. The following discussion is general, and will remain valid irrespective of any specific ordering of the locations.

2.2.1 Correlation matrices

All correlation matrices use daily data at the grid level and the correlations are calculated between grids. Depending on the situation, these matrices are calculated for the entire period (or for subsets of days, such as years, months within a year, etc.), and for the entire grid (or for subsets of grids such as climatic regions, etc.).

The correlations have been computed based on three different data matrices: $D = ((d_{dg})), T = ((t_{dg}))$ and $S = ((S_{dg})) d = 1, \ldots, 26298$, and $g = 1, \ldots, 280$, which represent respectively, the original DTR data, the residuals from time series decomposition

of the original data, and the trimmed DTR data (trimming is explained later).

Two different correlations have been used. One is the *product moment (Pearson's)* correlation, and the other is the *Bergsma's correlation*. The latter has been used in the context of spatial association measures. Specifically, correlation matrices that have been computed are listed below:

(1) Two Pearson correlation matrices R^D and R^S , based on D and S respectively.

(2) Two sets of Pearson correlation matrices $\{R_i^D\}$ and $\{R_i^S\}$ based respectively on the *i*-th year sub-matrices of D and S.

(3) Two sets of Bergsma's correlation matrices $\{B_i^D\}$ and $\{B_i^S\}$ based respectively on the *i*-th year sub-matrices of D and S.

(4) Bergsma's correlation matrices, $\{B_{i,k}^S\}$, based on S for each year *i*, and only for grids from each climatic region k.

(5) Bergsma's correlation matrix B_{mi}^S based on S for the *m*-th month of the *i*-th year and all grids.

2.2.2 Empirical spectral distribution; Marčenko-Pastur law

The sample correlation matrix is a real symmetric random matrix. The study of the eigenvalues of a random matrix occupies a central position in random matrix theory (RMT). A basic notion is the following: for any real symmetric $n \times n$ random matrix R_n , let $\lambda_1, \ldots, \lambda_n$ be its eigenvalues (which are all real). Then the empirical spectral distribution (ESD) of R_n is defined by

$$F_{R_n}(x) = \frac{\#\{i \le n : \lambda_i \le x\}}{n}.$$
 (2.1)

This corresponds to the probability distribution which puts mass n^{-1} at each of the eigenvalues. Roughly speaking this is the *histogram* of the eigenvalues of R_n . We shall use the ESD for comparing sample correlation matrices obtained using various contextual subsets of our data.

The Marčenko-Pastur (MP) law provides the asymptotic distribution of the ESD of the sample covariance matrix when dimensions grow and the observations are independent and identically distributed (iid, pre noise). In particular, suppose X is an $n \times p$ matrix with iid observations that have mean 0, variance 1 and finite fourth moment. Suppose $n \to \infty$ and $n^{-1}p \to y$, $0 < y < \infty$. Then the ESD of $n^{-1}X^{\top}X$ (the symbol \top denotes the transpose of a matrix) converges almost surely to what is called the MP law. For details, please refer to Bose [2021]. This law has a point mass $1 - y^{-1}$ whenever y > 1. Elsewhere (irrespective of the magnitude of y) it has the density

$$f(x) = \frac{1}{2\pi}\sqrt{(y_+ - x)(x - y_-)}, \quad y_\pm = (1 \pm \sqrt{y})^2, \quad y_- \le x \le y_+.$$

Note that this limit law does not depend on the underlying distribution of the variables. Indeed, it is also known that with a high probability, there are no stray eigenvalues to the left of y_{-} or and right of y_{+} . This can be used to distinguish signal from noise, by identifying the eigenvalues that are near the edges of the MP law. For example, if p = 280 and n = 365, then y = 0.767123, and $y_+ = (1 + \sqrt{y})^2 = 3.519$, $y_- = (1 - \sqrt{y})^2 = 0.0154$. For a random correlation matrix, the eigenvalues that lie within this range approximately may be attributed to noise, and the rest may be attributed to signal.

Using the MP law, the original matrix can be approximated by using the significant eigenvalues and corresponding eigenvectors, which would be the de-noised version of the original matrix.

2.2.3 Singular value decomposition (SVD)

Let X be an $n \times p$ matrix with real entries. Then one can choose two real matrices U and V of orders $n \times n$ and $p \times p$ respectively, and an $m \times n$ rectangular diagonal matrix $\Sigma = ((\sigma_{ij}))$ (so that $\sigma_{ij} = 0$ if $i \neq j$) with real entries, such that, the columns $\{u_i\}$ and $\{v_j\}$ of U and V are the (orthonormal) eigenvectors of XX^{\top} and $X^{\top}X$, and

$$UU^{\top} = I_{n \times n}, \quad VV^{\top} = I_{p \times p}, \text{ and } X = U\Sigma V^{\top},$$

where $I_{n \times n}$ and $I_{p \times p}$ are identity matrices of dimensions n and p respectively. This is the singular-value decomposition (SVD) of X. The elements $\sigma_{ii}, 1 \le i \le \min(n, p) = r$, of Σ can and shall be chosen to be in decreasing order of i. They are called the singular-values of X. It is then clear that

$$X = \sum_{i=1}^{r} \sigma_{ii} u_i v_i^{\top}.$$
(2.2)

This decomposition breaks up the matrix X into simpler components, and all the information about the matrix X are coded in the singular-values $\{\sigma_{ii}\}$ and the *left and right singular-vectors* $\{u_i\}$ and $\{v_j\}$. The matrices U and V induce *rotation* whereas the matrix Σ induces *scaling*. The SVD automatically adjusts for any change in the ordering of rows and columns of X.

The crucial idea is that the singular-values that are large in magnitude are more important, while those that are of lower magnitude could possibly be ignored. If we drop these smaller singular-values and consider a reduced sum on the right side of (2.2), then it will serve as a good low dimensional approximation for X, especially in situations such as ours where n is large. This also helps in computations, and moreover, when X is data dependent, as in our case, the retained singular-values and singular-vectors become important data analytic tools.

We use a trimmed version of the SVD in our analysis. We also use generalised SVD, which is a more sophisticated tool and is discussed briefly in Section S1.

2.2.4 Bergsma's correlation; Spatial Bergsma statistic

A value of 0 for the conventional product-moment correlation (i.e. Pearson correlation) does not imply independence. Many (non-parametric) alternatives try to address this and other issues. One of these is Bergsma's correlation ρ (see Bergsma [2005]), which

has the nice property that zero correlation implies independence. This is not a rank based method, like the Bergsma's τ correlation. For some details about its description, properties and estimation, see Bose et al. [2023]. We chose this correlation since it leads to a nice spatial measure of association that we shall discuss below.

To investigate spatial association, we need a statistical measure for it. Such a measure is built using two ingredients. Suppose we have p spatial units (for us, p = 280).

A spatial proximity weight matrix $W = ((w_{ij}))_{1 \le i,j \le p}$ represents the extent of proximity between the spatial locations. A numerical weight w_{ij} is assigned to each pair (i, j), and larger weights signify greater spatial proximity. See Getis and Ord [2010] for different popular choices for W. It is always assumed that $w_{ii} = 0$ for all *i*. Further, Wshall be row-standardized, so that each row sum in the matrix is equal to one. We shall use two choices for W. The first is the lag-1 adjacency matrix

$$w_{ij} = \begin{cases} 1 & \text{if regions } i \text{ and } j \text{ are adjacent locations,} \\ 0 & \text{otherwise.} \end{cases}$$

The second choice of W is obtained by applying exponential decay on a notion of distance between the locations. Here simple Euclidean distances between the grids have been considered to construct the base distance matrix.

The second ingredient is a similarity matrix. Suppose that we have a series of observations $X_i = \{x_{m,i}, m = 1, \ldots, T\}$ at each of the p locations, $i = 1, 2, \ldots, p$ for T time points. The similarity matrix Y is defined as $S := ((sim_{ij}))$, where sim_{ij} is some measure of similarity (dependence) between X_i and X_j . In the literature on spatial studies, often the chosen measure of similarity is taken to be the product moment correlation. However, our choice for $sim(i, j) = \rho(X_i, X_j)$ where $\rho(X_i, X_j)$ is Bergsma's correlation between the variables X_i and X_j , mentioned earlier.

The spatial association measure (Spatial Bergsma) S_B proposed by Kappara et al. [2023] uses Bergsma's correlation as follows. Let X_i denote a real variable for the location $i = 1, \ldots, p$. Then S_B with a row-standardized W is defined as

$$S_B = p^{-1} \sum_{1 \le i < j \le p} (w_{ij} + w_{ji}) \rho(X_i, X_j).$$
(2.3)

Suppose that we have observation vectors $X_i = \{x_{m,i}, m = 1, ..., T\}$. Let $\tilde{\rho}^{(ij)}$ denote the estimated ρ , calculated for the pair (X_i, X_j) (see Bose et al. [2023] for details). Then an estimate of S_B is the spatial Bergsma's statistics, given by

$$\tilde{S}_B := p^{-1} \sum_{1 \le i < j \le p} (w_{ij} + w_{ji}) \tilde{\rho}^{(ij)}, \qquad (2.4)$$

The **R** code in Bose et al. [2023] (Appendix) can be used for computing \tilde{S}_B . Computation of $\{\tilde{\rho}^{(ij)}\}$ involves computation of several *U*-statistics with complicated kernels. When the number of observations is large, this is computationally very intensive. Thus, we use this concept only on the yearly average data.

The asymptotic distribution of \tilde{S}_B , as the number of observations becomes large, is known only when the observations are independent and identically distributed across time. This may appear to be a major hurdle since the DTR values are not expected to show any sort of independence. Nevertheless, this can still be used fruitfully as a comparative and a diagnostic tool once the data is reasonably time-detrended as we have done in our analyses.

3 Preliminary exploration

Our primary focus of study is spatial association, and temporal changes therein, in the DTR data arising from India. As mentioned earlier, complete data is available only for 280 grid points out of 362, and most of our analysis is restricted to these. However, for certain statistics involving averages, we do use the partial data available for the additional 82 grids points.

3.1 Empirical spatial distributions of average DTR

We first focus on the empirical spatial distribution of average DTR over various temporal subsets. For this part we use data from all the 362 grids. The data can be grouped in various ways, for example, year-wise, or season-wise within each year. The locations can be stratified according to climatic regions too.



Figure 2: Yearly average DTR values (1951–2020) for the 362 grids.

In Figure 2, for each year on the horizontal axis, we have plotted the yearly average DTR values from the 362 locations on the vertical axis. Visually, the dispersion of the underlying average DTR distribution has generally decreased over the years, and there is visible shrinkage of distribution from both sides. There is also an indication of change in the pattern somewhere in the 1960's.

The yearly weather of India is generally divided into four seasons: December to February (DJF), March to May (MAM), June to August (JJA), and September to November (SON). In Figure 3, for each of the 72 years on the horizontal axis, we have



Figure 3: Seasonal average DTR values (1951–2020) for the 362 locations. In *clock-wise direction*; DJF (Dec.–Feb.), MAM (March–May), SON (Sept.–Nov), JJA (June–Aug.).

plotted the DTR averages of each season (*seasonal averages*) for the 372 grids on the vertical axis. The behaviour of the DTR varies across seasons. The shrinkage of the distribution that was visible in Figure 2 is particularly pronounced during autumn (SON), and to a lesser extent in winter (DJF). The spatial spread of DTR is narrowest and most homogeneous during monsoon (JJA) and is widely varying in the other three seasons.

The average DTR values could also be plotted after stratifying the grids by climatic zones or by climatic zones and seasons. Some such plots are given in Section S1.

3.2 Empirical spectral distribution of DTR correlation matrices

We now consider the (product-moment) correlation matrices $\{R_i^D\}$ obtained from the daily DTR values for all pairs of the 280 grid points, for every year *i*. We shall study these random matrices to assess the changes in the underlying spatial association pattern over the years, as has been suggested by Figures 2 and 3.

In Figure 4a we have presented the ESD for each year, based on R_i^D . Over the years there has been a small but steady increase in the median value of the ESDs, and a slight shift away from zero in the left tail of the distribution. Since the sum of eigenvalues is fixed for a correlation matrix, the above, generally speaking, will mean that the right



Figure 4: ESD analyses of the 72 correlation matrices $\{R_i^D\}$.

tail shrinks towards the middle too. In Figure 4b, we have plotted the quantiles (at 10% interval) of the eigenvalues for each year. We notice a small but gradual decrease in the top eigenvalues, although noticeable oscillation in these values is also present.

3.3 Correlation matrix R^D and its de-noised version \hat{R}^D

Let $\{\lambda_i^D\}$ be the eigenvalues of R^D . Applying the MP law based cutoff, only 10 of these were significant. Then R^D is approximated by using these eigenvalues and the corresponding eigenvectors $\{e_i^D\}$ as the the MP law based *de-noised matrix* \hat{R}^D :

$$\hat{R}^D = \sum_{i=j}^{10} \lambda_j^D e_j^D (e_j^D)^\top.$$

The upper and lower triangles in Figure 5 are based on R^D and \hat{R}^D respectively. First note that from the upper triangle, there is visible association between spatial locations. Indeed, of the 10 top eigenvalues, the largest is ten times in magnitude compared to the second largest, as is evident from Figure 4b. This may not be surprising since common temporal pattern arising due to seasonality would naturally lead to a strong association/correlation. Nevertheless, presence of high degree of positive association might be masking other patterns. MP law based de-noising has some effect and brings balance to the distribution of the correlations (observe the negative values along the left vertical). However, the correlation distribution still looks heavily positive.

Thus although the MP law has been very useful in many applications, it is not effective here in bringing out any pattern, when implemented on R^D .



Figure 5: Upper triangle-entries from the correlation matrix R^D ; Lower trianglecorresponding entries from the MP de-noised version of R^D .

4 Analysis

The preceding explorations indicate that in order to bring out meaningful spatial association patterns, it is imperative to remove association due to predominant temporal patterns. Time series decomposition was carried out for each of the 280 daily-DTR-time series. Time trends and seasonal components were removed and the spatial correlation matrix R^T was obtained based on the resulting time-detrended data T. By applying MP law, there were 18 significant eigenvalues and once again there is only one dominant eigenvalue. Thus we explore alternate methods for removing dominant trends to bring out the spatial association pattern underlying the DTR data.

4.1 Trimmed-DTR matrix S: detrending by SVD

SVD is used to project the data onto a lower ranked subspace. The common strategy is to use the truncated SVD after dropping the relatively smaller singular-values.

However the objective here is to remove the predominant patterns to elucidate the finer and yet unknown (spatial) patterns. For D, the presence of a few (predominantly temporal) strong signals are masking other (spatially) significant signals. Our strategy would be to remove components corresponding to a few dominant singular-values of D, and work with the resulting *trimmed DTR matrix S*.

For that purpose we identified the minimal number of singular-values D required to reasonably *time-detrend* the data, while retaining as much of the overall signal as possible. Efficacy of time detrending was assessed by computing the *autocorrelation* function (ACF), and the overall signal coverage was measured by the *percent share of*



Figure 6: Cumulative singular-values for the original DTR data D.



Figure 7: Autocorrelation series for the 280 time series based on trimmed data S.

singular-values retained. Employing these two criteria, we stopped after the top 12 singular-values. Referring to Figure 6, these cover 72% of the singular-values. Taking out the components $\sigma_{ii}u_iv_i^{\top}$ from $D, 1 \leq i \leq 12$, we obtain the trimmed data matrix

$$S = D - \sum_{i=1}^{12} \sigma_{ii} u_i v_i^{\top}.$$
 (4.1)

The autocorrelation plots of the time series from each row of S given in Figure 7 show that this trimming removed a lot of temporal patterns.

Now consider the correlation matrix \mathbb{R}^S obtained from S. Applying the MP law on the ESD of \mathbb{R}^S now shows 33 significant eigenvalues, compared to only 10 for \mathbb{R}^D (and 18 for the correlation matrix based on the residuals from time-series decomposition T, as mentioned earlier). See Figure S6 on eigenvalues of correlation matrices in Section S1.

In addition to these, we also consider the entire ESDs. Figure 8 shows that the ESD of R^D is extremely long tailed. The ESD for R^S has a much more balanced



Figure 8: Eigenvalues of the correlation matrices: top panel, R^D ; bottom panel, R^S .

distribution. Trimming has shifted the spatial association behaviour, as captured by the yearly correlation matrices, and also made the distributions more balanced. Later in Section 4.4 we have also considered the behaviour of the spectral distribution of the yearly correlation matrices $\{R_i^S\}$. In Section S1 we have given further information on the singular-vectors removed.

4.2 Generalised SVD to assess spatio-temporal information in S

The method employed above is not a traditional de-trending technique. We need to assess if there has been loss of spatial information. We use GSVD for this purpose. See Section S1.5 for a brief description of GSVD. For each successive pairs of years data from both D and S (resulting in 71 such pairs from each), we computed the GSVD. The results are shown in Figure 9, where we have given the boxplots for the log-generalised singular-values for all such successive pairs, separately for D and S.

Recall that in GSVD the right singular-vectors are held common while the left singular-vectors are allowed to vary between the two matrices. In the current context this translates into holding the vectors capturing the spatial rotation fixed, while allowing the temporal rotation to vary from year to year.

As Figure 9 shows, the empirical distributions of the generalised singular-values remain same or similar for pairs arising from D and S. This confirms that the temporal information retained in S is comparable to that in D.

Recall that to identify significant signals, the MP law was used for the eigenvalues of the correlation matrices. To carry out similar exercise for the generalised singular values, we needed a similar result on the null distribution of generalised singular values. This seems to be missing in the literature.



Figure 9: Empirical distributions of the (log) generalised singular-values based on pairwise successive years data from D (left panel) and S (right panel). The (empirical) critical values for the generalised SVs are marked by horizontal lines.

Hence we estimated the null distribution of GSVs empirically (see Section S1.6 and Figure S1) and obtained the critical values from that. By comparing the observed GSV distributions and the critical values, it appears that (in the log scale) the GSVs for S has a more "regular" distributional behaviour compared to that for the original DTR matrix.

We have also carried out the reverse exercise to assess spatial information retention using D^{\top} and S^{\top} (see Section S1.7) with similar conclusion as above.

This exercise assures that S can be used as proxy for "time-detrended" series, and it retains the spatial features of D. Hence S is suitable for studying potential spatial patterns in the DTR data, and will be considered next for further analyses.

4.3 Exploring spatial association pattern

Now we consider the correlation matrix R^S obtained from S. Prior to ascertaining existence of any pattern within that matrix, we first explored it from various perspectives.

For example we identified a few selective grids that cover some of the major cities in India, and plotted the correlation vectors in the form of heat maps corresponding to the latitude and longitude information of individual grids. The resulting images given in Figure 10 exhibit strong spatial patterns around these cities. There appears to be strong association in the immediate vicinity for each of these cities.

Let us fix a grid g_1 , and consider the grid g_2 which yields the maximum correlation with g_1 . Consider the difference in latitudes and longitudes between g_1 and g_2 . The distribution of these differences across all grids has a significant mode at 0, while the



Figure 10: Correlations within grids for S that cover each of the four major cities. top left, Delhi; top right, Kolkata; bottom left, Chennai; bottom right, Mumbai.



Figure 11: Distribution of distance between grids with maximum association, measured in terms of differences in latitude (in black) and longitude (in red) for S.

distribution of the difference in longitudes is approximately uniform around ± 1 . In Figure 11) we have plotted these two distributions. Such salient patterns were not visible earlier but are emerging now. The question remains whether we can present this pattern for all grids via a single plot. We first stratify the grids according to the climatic regions they belong to. Following which we employ the Hilbert space filling curve approach and arrange the grids in a spiral manner described earlier. We present in Figure 12 the correlation matrix with the grids ordered as described above. This enables us to visualize much more clearly the spatial association pattern, and also indicates variation in strength and range where maximum association between locations take place. Analysis of the eigenvectors corresponding to significant eigenvalues as per the MP law support the observation of climatic region-specific spatial association pattern.

4.4 Exploring temporal changes in spatial association pattern

The inherent assumption in obtaining the preceding pattern is that it is temporally static, and hence the entire data was used to arrive at the conclusions. To study similar association patterns for any smaller time interval, for example individual calendar year, would result in the task of summarising numerous such matrices.

We achieve this by studying the ESD of each of the correlation matrices $\{R_i^S\}$. Studying the ESD in lieu of the entire correlation matrix enables us to investigate for patterns in a much smaller dimension while retaining the essential features of the data. Figure 13 presents the ESD for each year. The ESDs clearly indicate the occurrence of a drastic change in the spatial association pattern in the late 1960s.

4.5 Spatial Bergsma: A univariate spatial association measure

The correlation matrices present a 280×280 dimensional summary of the association, and the ESDs present a 280×1 dimensional summary. We use spatial Bergsma measure of Kappara et al. [2023] to arrive at a univariate measure for the spatial pattern observed.



Figure 12: Correlations based on S with grids arranged first according to climatic regions and then in spiral Hilbert space filling curve manner. Upper triangle, original correlation matrix; Lower triangle, the MP de-noised version.





Figure 14: Spatial Bergsma measure with two spatial weight matrices for S: *red*, lag-1 adjacency; *blue*, exponential distance decay.

As described earlier in Section 2.2.4, the key ingredients to derive such a measure are a spatial weight matrix and a similarity measure. We used two candidates for spatial weight matrix, namely the lag-1 adjacency matrix and the exponentially decaying distance matrix. For similarity measure we would be using Bergsma's ρ for its useful ability to depict independence (corresponding to the zero value of the measure).

Figure 14 presents the S_B statistics at the all India level for each year based on S along with the same for each of the six climatic regions. It illustrates that both degree of association and variation in association, are climatic region specific. The S_B statistic shows without any doubt that there has been at least one, and possibly more, major change(s) in spatial association pattern over the Indian region. Such changes in association has had varying degrees of effect in different climatic regions.

4.6 El Niño Southern Oscillation and spatial association

Figures 12 and 14 have already established that climatic regions have different impacts on the centrality and the dispersion of the spatial association measure S_B for India. It is well



Figure 15: Boxplot of spatial Bergsma statistics, based on individual year's S and categorized according to ENSO information.

known that El Niño-Southern Oscillation (ENSO) is a key global weather phenomenon, and affects the climate of much of the tropics and subtropics (see https://www.noaa.gov/for more information).

We augmented the yearly S_B statistics values obtained in the earlier section with the ENSO data. Figure 15 shows the impact of ENSO on the distributional behaviour of S_B . We can see that the variations in S_B are much less in the El Niño phase, and it can be argued that this is because the influence of the phase is so strong that it takes over and dictates the variation in DTR.

5 Conclusions

Data explorations indicate that in order to bring out meaningful spatial association patterns, it is imperative to remove association due to predominant temporal patterns. The trimmed DTR matrux S, trimmed-off by the top 12 singular-values, and the corresponding singular-vectors, was found suitable as proxy for "time-detrended" series, with significant spatial signals still contained in it. Unsurprisingly, there is a strong association in the immediate vicinity for each of the grids, with patterns slightly differing in north-south and east-west directions. However such association should not be taken as indicative of causation.

Overall, the ESD based analyses support the presence of a climatic region-specific spatial association pattern. Climatic regions have had different impacts on the centrality and the dispersion of a spatial association measure. We can see that associations have also been affected by global climatic phenomenon like ENSO, in particular by the El Niño phase. Influence of global factors have been reported in the literature earlier also Vinnarasi and Dhanya [2019].

The yearly ESD based analyses clearly indicate a drastic change in the spatial association pattern in the late 1960s. The S_B statistic confirms without any doubt that there has been at least one, and possibly more, major change(s) in the spatial association

pattern in the Indian region.

While most of the methodological techniques used here exist in the literature, a novel synergistic manner of employing these techniques to data can (and did) bring out new features. In the current situation it brought out yet unknown spatial association patterns and temporal breakage in it. Such a pipeline of analysis can be deployed to other investigative studies on spatio-temporal data as well.

We have obtained the null distribution of generalised singular values empirically, and have used it profitably to assess the eigenvalues that have significant signals. This is due to the fact that distributional results on the generalised singular values for random data matrices appear to be absent in the literature. It is worth pursuing this from the theoretical perspective.

We are in broad agreement with the findings of other researchers that DTR showed an increase during early seventies till the turn of the millennium (see for instance Mall et al. [2021]). This is clear even while accounting for spatial variation and using a more robust measure (than mean) against outliers.

In future the spatial association/dependence pattern that emerged in our analysis could help build composite spatio-temporal models, e.g. by using separable Gaussian space-time processes.

References

- W. Bergsma. On a new correlation coefficient, its orthogonal decomoposition, and associated tests of independence. Available at SSRN 905093, 2005.
- A. Bose. Random Matrices and Non-Commutative Probability. CRC/Chapman-Hall, 2021. doi: 10.1201/9781003144496.
- A. Bose, D. Kappara, and M. Bhattacharjee. Estimating Bergsma's covariance. Jour. Korean Stat. Soc., 2023. doi: 10.1007/s42952-023-00236-1.
- P. M. de F. Forster and S. Solomon. Observations of a "weekend effect" in diurnal temperature range. *Proc. of Nat. Acad.of Sciences*, USA, 100(20):11225–11230, 2003. doi: 10.1073/pnas.2034034100.
- A. Getis and J. K. Ord. The analysis of spatial association by use of distance statistics. In *Perspectives on Spatial Data Analysis*, pages 127–145. Springer, 2010.
- D. Jhajharia and V. P. Singh. Trends in temperature, diurnal temperature range and sunshine duration in Northeast India. Int. J. Climatol., 31:1353–1367, 2011.
- F. Ji, Z. Wu, J. Huang, and E. P. Chassignet. Evolution of land surface air temperature trend. Nat. Clim. Chang., 4:462–466, 2014.
- D. Kappara, A. Bose, and M. Bhattacharjee. An association measure for spatio-temporal time series. *Metrika*, 2023. doi: 10.1007/s00184-023-00939-9. Special Issue on Categorical Data. Print version to appear.

- W. Köppen. Klassification der klimate nach temperatur, niederschlag and jahreslauf. Petermanns Geographische Mitteilungen, 64:193–203, 243–248, 1918.
- D. R. Kothawale, J. V. Revadekar, and R. K. Kumar. Recent trends in pre-monsoon daily temperature extremes over india. J. Earth Syst. Sci., 119:51–65, 2010.
- R. K. Mall, M. Chaturvedi, N. Singh, R. Bhatla, R. S. Singh, A. Gupta, and D. Niyogi. Evidence of asymmetric change in diurnal temperature range in recent decades over different agro-climatic zones of India. *Int. Jour. Climatology*, 41(4):2597–2610, 2021. doi: 10.1002/joc.6978.
- T. L. McKnight and D. Hess. Climate Zones and Types. Physical Geography: A Landscape Appreciation. Prentice Hall, New Jersey, 2000.
- S. Sen-Roy. Spatial patterns of trends in seasonal extreme temperatures in India during 1980–2010. Weather and Climate Extremes, 24:100203, 2019. doi: 10.1016/j.wace. 2019.100203.
- A. Sharma, D. Swami, and N. Joshi. Analysis of temperature extremes over India using Sheffield & IMD dataset. *Research Square*, 2021. doi: 10.21203/rs.3.rs-662326/v1. Preprint.
- C. W. Stjern. How aerosols and greenhouse gases influence the diurnal temperature range. Atmos. Chem. Phys., 20:13467–13480, 2020. doi: 10.5194/acp-20-13467-2020.
- R. Vinnarasi and C. T. Dhanya. Quantifying the shifts and intensification in the annual cycles of diurnal temperature extremes for human comfort and crop production. *Environ. Res. Lett.*, 14:054016, 2019. doi: 10.1088/1748-9326/ab0fe5.
- R. Vinnarasi, C. T. Dhanya, C. A., and A. AghaKouchak. Unravelling diurnal asymmetry of surface temperature in different climate zones. *Sci Rep.*, 7(1), 2017. doi: 10.1038/s41598-017-07627-5.
- Z. Wu, N. E. Huang, and X. Chen. The multi-dimensional ensemble empirical mode decomposition method. Adv. Adapt. Data Anal., 01:339–372, 2009.
- Z. Zhong, B. He, H. Chen, D. Chen, T. Zhou, W. Dong, C. Xiao, S.-P. Xie, X. Song, L. Guo, R. Ding, L. Zhang, L. Huang, W. Yuan, X. M. Hao, D. Ji, and X. Zhao. Reversed asymmetric warming of sub-diurnal temperature over land during recent decades. *Nature Comm.*, 14:7189, 11 2023. doi: 10.1038/s41467-023-43007-6.

S1 Supplementary information

We probe the DTR values further, using stratifications via climatic zones and seasons. The six climatic zones will be numbered as: (1) Tropical monsoon, (2) Tropical savannah, wet and dry (3) Arid, Steppe, hot, (4) Humid subtropical, (5) Montane climate, and (6) Hot deserts, Arid.

S1.1 Exploration of average DTR values by climatic zones and seasons

As mentioned earlier, we could use stratification by climatic zones, and by seasons within each climatic zone.



Figure S1: Average DTR values for 72 years at the 362 locations in the six climatic zones, the six zones (1-6) presented in the order left to right and top to bottom panel.

Figure S1. Yearly average DTR, in each climatic zone: As mentioned earlier, India has six climatic zones. DTR behaviour across these zones are quite different. Nevertheless, the change in pattern somewhere in the 1960's observed earlier in the aggregated yearly data is also visible in each of these zones. There appears to be a systematic oscillation in the average DTR values for the desert climatic region. It is quite pronounced when checked at the four seasons level (approximately every 5-6 years).



Figure S2: Average DTR values for Dec. to Feb. for 72 years, at the 362 locations in the six climatic zones (1)–(6) in the order left to right and top to bottom panel.

Figure S2. Yearly average DTR, in each climatic zone and season: We can also plotting the DTR using both the seasons and the zones. This is done for the winter season (December-January-February) in this figure. Sudden change in location of the DTR distribution is seen even at a climatic region level for any given season.

S1.2 Empirical singular-value (SV) distribution

The significant top singular-values of D could be identified if the null distribution of singular-values for a random matrix of the same size were available. The limit distributions of such high-dimensional matrices as the dimensions grow, are known and are

universal, in that they do not depend on the underlying distribution of the standardized i.i.d. entries and could be used as the null distribution. Instead, we obtained the empirical distribution of the singular-values based on 1000 simulations of matrices of order 26298×280 with N(0,1) entries. Figure S3 provides this empirical distribution. Based on the scaled DTR data, at 5% (two-sided) significance levels, there were only 3 singular-values out of 280, that were not found significant.



Figure S3: Empirically estimated null distribution of singular-values.

S1.3 Further findings from SVD of D

Figures S4 and S5 present some aspects of the top 6 left singular-vectors of D. Unsurprisingly they have association with the seasonal patterns as captured at the month level. It has been reported in the literature (for example, de F. Forster and Solomon [2003]) that weekdays and weekends can have different impacts on DTR. Figure S5 also suggests that there could be an association between these variates.

S1.4 Eigenvalues of R^D , R^T and R^S

The spatial association of DTR over the Indian subcontinent could be studied via the ESDs from the correlation matrices based on D, T or S. Figure S6, clearly shows presence of a handful of dominant eigenvalues for the first two, and a more reasonable decay of eigenvalues for the third.

S1.5 Generalised singular-value decomposition (GSVD)

GSVD is a simultaneous linear transformation of two matrices D_1 and D_2 of orders $N_1 \times M$ and $N_2 \times M$. In implementations it is assumed that the number of rows should be more than the number of columns, so that $\min(N_1, N_2) \ge M$.

Before we give a description of GSVD let us indicate what D_1 and D_2 shall be in our context. Let the matrix D_1 be the $N_1 \times M$ sub-matrix of D whose *n*-th row vector



Figure S4: Components of top left singular-vectors of D, at monthly level.



Figure S5: Components of top left singular-vectors of D, at weekday/weekend level for each year.



Figure S6: Eigenvalues of correlation matrices R^D , R^T and R^S .

holds the DTR values for the *n*-th time point across all the M = 280 grids. The *m*-th column of D_1 holds the DTR values of the *m*-th grid across different time points. Let D_2 be the $N_2 \times M$ corresponding matrix for the *subsequent year*. This one-to-one correspondence between the two sets of conditions is at the foundation of the GSVD comparative analysis of the two data sets.

Here we use GSVD to obtain a simultaneous linear transformation of the two matrices D_1 and D_2 to reduced $M \times M$ space, with shared row-space basis vectors as follows:

$$D_1 = U_1 \Delta_1 [0 \ P] V$$
 and $D_2 = U_2 \Delta_2 [0 \ P] V$

where U_1 , U_2 and V are orthogonal matrices of dimensions $N_1 \times N_1$, $N_2 \times N_2$ and $M \times M$ respectively.

Let r be the rank of the matrix $\begin{bmatrix} D_1^\top & D_2^\top \end{bmatrix}$, with $r \leq M$. Then P is an $r \times r$ upper triangular non-singular matrix, and $\begin{bmatrix} 0 & P \end{bmatrix}$ is an $r \times M$ matrix. In the case of DTR data, there is no additional information available to assume that the matrices may be rank deficient and thus it is expected that r = M.

The matrices Δ_1 , Δ_2 are non-negative quasi diagonal matrices of orders $N_1 \times r$ and $N_2 \times r$ respectively, and further,

$$\Delta_1^{\top} \Delta_1 + \Delta_2^{\top} \Delta_2 = I,$$

where I is the $r \times r$ identity matrix.

Write

$$\Delta_1^{\top} \Delta_1 = \operatorname{diag}(\alpha_1^2, \dots, \alpha_r^2) \text{ and } \Delta_2^{\top} \Delta_2 = \operatorname{diag}(\beta_1^2, \dots, \beta_r^2),$$

where $0 \leq \alpha_i, \beta_i \leq 1$ for all *i*. The ratios $\alpha_1/\beta_1, \ldots, \alpha_r/\beta_r$ are called the generalized singular values of the pair (D_1, D_2) . If $\beta_i = 0$, then the generalized singular value α_i/β_i is infinite.

The GSVD is unique up to phase factors of ± 1 of each triplet of corresponding column and row basis vectors, except in degenerate subspaces defined by subsets of pairs of generalized singular values of equal ratios. Note that the basis vectors for row spaces

of the two matrices are shared (i.e. held common between the two matrices). Thus these generalised singular values capture useful essence of the two matrices relative behaviour with regard to the column spaces.

The anti-symmetric angular distances between D_1 and D_2 are,

$$\theta_i = \arctan(\alpha_i/\beta_i) - \pi/4, 1 \le i \le r.$$

They indicate the relative significance of the *i*-th basis vectors from the column spaces of D_1 and D_2 , and are arranged in decreasing order of significance in D_1 relative to D_2 .

S1.6 Empirical generalised singular-value (GSV) distribution

We were unable to locate any theoretical results on distributional behaviour of the generalised singular values in literature. We propose to use Empirical GSV distribution instead. We implemented two alternate methods to simulate the null situation. In the first method, we simulate data of size similar to our current analysis setup and derive GSVs. We repeat this process 1000 times and use the collected GSVs to form an empirical distribution. In the second method, we use (random ??) permutations on the actual data and obtain the GSVs in each case, repeating the process 1000 times, to obtain a second empirical distribution. There were no discernible differences in these two estimates of the empirical distribution of GSVs.



Figure S7: Simulated null distribution for 366×280 and 365×280 matrices with independent N(0,1) elements, based on 1000 simulations: *left panel*-GSV; *right panel*-log GSV); *top panel*-boxplots; *bottom panel*-estimated overall density.



Figure S8: Empirical distributions of the (log) generalised singular-values based on transposed half-yearly data (i.e. in grid x day organisation), with original DTR data in left panel and SVD-based-trimmed DTR data in the right panel. The (empirical) critical values for the generalised SVs are marked by horizontal lines.

S1.7 GSVD on DTR data and spatial information

To assess retention of spatial information in S compared to D, we could consider the sub-matrices pertaining to each year's data from D^{\top} and S^{\top} . In that case the number of rows will be 280 from the 280 spatial locations/grids and number of columns would be 365 or 366.

However, a technical requirement in the computation of GSVD is that the number of columns be less than the number of rows for the pair of matrices being decomposed. Thus we can not simply use the transposed matrices D_1^{\top} and D_2^{\top} and proceed as before. So we consider each year's data matrix D_i and S_i and further split it into two equal sized matrices with row sizes 182 (or 183) corresponding to the two halves of a calendar year, ignoring a singleton row for non-leap-years and consider the transpose these matrices. The results for GSVs on these "transpose matrices" are given in Figure S8. The end conclusion remains as before (see Figure 9), Figure S8 also shows that the empirical distributions of the generalised singular-values remain same or similar for the two matrices D and S in the spatial direction as well.

S1.8 Analysis of alternate DTR data for India

We employed the same analysis pipeline on an alternate data source for DTR for the Indian region from http://www.cru.uea.ac.uk/cru/data. We gathered monthly DTR data for the same 72 years and same 1 degree resolution in each of east-west and north-south direction. There were 258 grids covering India with complete data for all 72

years. Figure S9 shows the same spatial pattern from this monthly DTR data as was



Figure S9: Correlation matrices based on monthly DTR data of India from CRU. *left* panel–original DTR data; *right panel*–trimmed DTR with 12 singular values and grids arranged according to climatic regions).

identified by the daily DTR data. This provides substantial confidence in the novel analysis pipeline proposed here.

S1.9 Analysis of other geographic location: Bahia-Brazil

Daily DTR data from Bahia of Brazil for the period January 01, 2000 to December 31, 2022 was considered. This data is not gridded, unlike the data from India considered in this manuscript. Instead, each of the 417 municipalities in Bahia contributes to an individual time series. Overall coverage is approximately about 10 degrees in both east-west and north-south direction. Thus overall this data covers a relatively smaller area with higher number of time series. Figure S10 elucidates novel spatial pattern for



Figure S10: Correlation matrices based on daily DTR data from 417 municipalities of Bahia, Brazil. *left panel*—original DTR data; *right panel*—trimmed DTR with 5 singular values with municipalities arranged according to spatial proximity).

Bahia, although the underlying data is neither gridded, nor at similar spatial resolution. Also temporal coverage in this data is much shorter (only 23 years) and climatological coverage is also limited, since Bahia has mostly tropical wet and dry or savanna climate (which is one of 6 climatic regions in India).

S2 Notes on figures

We present below, at one place, a description of all the figures given in this article. "DTR" data is primarily in a matrix form $((X_{ij}))_{i \in t, j \in s}$, where t is an index set for time, which could represent day/season/year, and s is an index set for space, which could be the grids of size $1^{\circ} \times 1^{\circ}$ that cover India, or a climatic subset of those. The "DTR" data could be the original DTR data D, the time series decomposition based detrended DTR data T, or the SVD based detrended trimmed DTR data S.

For original daily DTR data, D, t spans 26298 days covering 1-Jan-1951 to 31-Dec-2022 and s covers the 362 grids that cover India. We have complete data for the 72 years (1951-2022) for 280 of these grids, so often the size of s is 280. However s could also represent one of the six climatic zones of India.

In various figures DTR values have been summarised using functions like average over a (sub)set of time and/or space, correlation matrices, empirical spectral distributions, singular-values, singular-vectors, eigenvalues, generalised singular-values, etc. We use the following notation to describe the figures:

 $D = ((d_{ij}))_{26298 \times 280}$ is the data matrix.

 $U\Sigma V^{\top} = D$ is the singular-value decomposition, where $\Sigma = Diag(\sigma_j, j = 1, ..., 280)$, and the singular-values $\{\sigma_j\}$'s are arranged in decreasing order.

 $S = ((S_{ij}))_{26298 \times 280} = D - \sum_{j=1}^{12} \sigma_j u_j v_j^{\top}$ is the trimmed data matrix.

 $T = ((S_{ij}))_{26298 \times 280}$ is the residual matrix after a time series decomposition of each column of D using additive models each consisting of a trend component (estimated by simple moving average), a seasonal component and a random/residual component.

 R^D is the 280 × 280 correlation matrix based on D.

 R^D_i is the 280×280 correlation matrix based on D for the i-th year.

 R^S is the 280 × 280 correlation matrix based on S.

 R_i^S is the 280 × 280 correlation matrix based on S for the *i*-th year.

 $\{\lambda_j^D\}$ and $\{\lambda_j^S\}$ are the eigenvalues, in decreasing order, of the correlation matrices R^D and R^S .

 $\hat{R}^D := \sum_{j=1}^{10} \lambda_j^D e_j^D (e_j^D)^\top \text{ is the de-noised approximation of } R^D.$ $\hat{R}^S := \sum_{j=1}^{33} \lambda_j^S e_j^S (e_j^S) \top \text{ is the de-noised approximation of } R^S.$

 SB_i^S : Spatial Bergsma statistic using S, for the *i*-th year.

 SB_{ij}^S : Spatial Bergsma statistic using S, for the *i*-th year in the *j*-th climatic zone.

Figure 1: India divided into $362 \ 1^{\circ} \times 1^{\circ}$ grids, each grid labelled by one of the six climatic region codes.

Figure 2: Boxplots for each of the 72 years, of the yearly averages from the 362 grids.

Figure 3: Boxplots for averages over the daily data from D, for each given season (DJF,

MAM, JJA, SON), year (i = 1, ..., 72) and grid (j = 1, ..., 362).

Figure 4: For every year *i*, the left panel are the bulk eigenvalues (presented in the form of boxplots without outliers) of R_i^D , The right panel is the plot of percentiles, 0(10)100, of these eigenvalues.

Figure 5. Heat map: the upper triangle is based on the correlations from R^D ; the lower triangle is for the corresponding entries from the matrix \hat{R}^D .

Figure 6: Plot of the cumulative sums of the singular-values of D.

Figure 7: 280 auto-correlation series for each column of S treated as a time series.

Figure 8: Box plots of the eigenvalues $\{\lambda_j^D\}$ and $\{\lambda_j^S\}$ of \mathbb{R}^D and \mathbb{R}^S .

Figure 9: Consider the daily trimmed DTR matrix S. Let S^{y_1} and S^{y_2} be data on two successive years $(y_1 \text{ and } y_2)$ from S. GSVD enables simultaneous decomposition of these matrices as

$$S^{y_1} = U_1 \Delta_1[0 \ P]V$$
 and $S^{y_2} = U_2 \Delta_2[0 \ P]V$

where U_1 , U_2 and V are orthogonal matrices, Δ_1 , Δ_2 are non-negative quasi diagonal matrices which satisfy $\Delta_1^{\top} \Delta_1 + \Delta_2^{\top} \Delta_2 = I$ and P is an upper triangular non-singular matrix. Then generalised singular-values (gsv) are the ratios of the respective elements from the diagonals of Δ_1 and Δ_2 . Typical implementation of this algorithm ensures these ratios are positive and decreasing. The figure gives 71 boxplots of these gsvs for each successive pairs of years. The figure also includes similar boxplots for the gsvs based on D.

Figure 10: For each of the four major cities, Delhi, Kolkata, Chennai and Mumbai, the correlations (from R^S), with the other cities have been plotted as heat maps.

Figure 11: Fix a grid g_1 , and consider the grid g_2 which yields the maximum correlation (based on \mathbb{R}^S) with g_1 . Consider the difference in latitudes and longitudes between g_1 and g_2 . The two distribution of these differences across all grids have been plotted.

Figures 13–12 are similar to Figures 4a–5 but based on S instead of D.

Figure 12: Heat map: the upper triangle is for the correlations from R^S ; the lower triangle is for the corresponding entries from the matrix \hat{R}^S . For both matrices, the grids are arranged first by climatic regions, and within each region in a spiral manner according to a Hilbert space filling curves.

Figure 13: For every year *i*, plot of bulk of the eigenvalues (as boxplot without outliers) of R_i^S .

Figure 14: Spatial Bergsma statistics S_B with two spatial weight matrices, W_1 (lag-1

adjacency, in red) and W_2 (exponential distance decay in blue), for S. In the left panel, the yearly SB_i^S , i = 1, ..., 72. In the right panel, the yearly SB_{ij}^S , for $1 \le i \le 72, 1 \le j \le 6$.

Figure 15: Boxplots of $\{SB_i^S\}$, i = 1, ..., 72 for the three ENSO categories.

Figure S1: Average DTR values for each of the 72 years at the 362 locations in the six climatic zones, the six zones (1-6) presented in the order left to right and top to bottom panel.

Figure S2: Average DTR values in the winter season (December to February) for each of the 72 years, at the 362 locations in the six climatic zones, with the six zones (1-6) presented in the order left to right and top to bottom panel.

Figure S3: Simulated null distribution of Singular Values based on matrices of dimensions 26298×280 with independent N(0, 1) elements and 1000 replications.

Figure S4: Components of the top 6 left singular-vectors of D, split according to months, yielding 6×12 boxplots.

Figure S5: Components of the top 6 left singular-vectors of D, split according to year, and within each year by weekday or weekend, yielding 72×2 plots for each component.

Figure S6: Eigenvalues of correlation matrices, based on D, S and T.

Figure S7: Simulated null distribution of Generalised Singular Values (*left panel*) and log GSV (*right panel*), with boxplots from each of 1000 simulations on top panel and estimated overall density in the bottom panel. Based on matrices of dimensions 366×280 and 365×280 with independent N(0, 1) entries.

Figure S8: For each year *i*, split S_i^{\top} into two approximately equal halves and perform GSV decomposition on the resulting two 280 x 182 (or 183) dimensional matrices. Figure shows boxplots of these GSV's from each year. Also included are similar boxplots based on D_i^{\top} .

Figure S9: This figure is similar to Figures 5 and 12 but with alternate data source and is based on monthly data instead of daily DTR data from India (and without MP de-noising).

Figure S10: This figure is similar to Figures 5 and 12 but with data from Bahia, Brazil (and without MP de-noising).