Computers & Geosciences 90 (2016) 78–86

Contents lists available at ScienceDirect

Computers & Geosciences

journal homepage: www.elsevier.com/locate/cageo

Research paper

Finding a needle by removing the haystack: A spatio-temporal normalization method for geophysical data

E. Pavlidou*, M. van der Meijde, H. van der Werff, C. Hecker

Department of Earth Systems' Analysis, ITC, University of Twente, Enschede, The Netherlands

ARTICLE INFO

Article history:
Received 3 September 2015
Received in revised form
27 January 2016
Accepted 24 February 2016
Available online 26 February 2016

Keywords:
Time series
Anomaly detection
Kernel-based approach
Near-real time
Satellite imagery

ABSTRACT

We introduce a normalization algorithm which highlights short-term, localized, non-periodic fluctuations in hyper-temporal satellite data by dividing each pixel by the mean value of its spatial neighbourhood set. In this way we suppress signal patterns that are common in the central and surrounding pixels, utilizing both spatial and temporal information at different scales. We test the method on two subsets of a hyper-temporal thermal infra-red (TIR) dataset. Both subsets are acquired from the SEVIRI instrument onboard the Meteosat-9 geostationary satellite; they cover areas with different spatio-temporal TIR variability. We impose artificial fluctuations on the original data and apply a window-based technique to retrieve them from the normalized time series. We show that localized short-term fluctuations as low as 2 K, which were obscured by large-scale variable patterns, can be retrieved in the normalized time series. Sensitivity of retrieval is determined by the intrinsic variability of the normalized TIR signal and by the amount of missing values in the dataset. Finally, we compare our approach with widely used techniques of statistical and spectral analysis and we discuss the improvements introduced by our method.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Short-term, localized, non-periodic fluctuations in hyper-temporal measurements are often obscured by background patterns in the data. The terms 'short-term' and 'localized' respectively refer to duration and spatial extent considerably smaller than the rest of the dataset. Such fluctuations are often of interest for geoscience applications based on detection of extremes and/or environmental monitoring. Potential examples include fires, volcanic and geothermal activity, fluctuations of climatic variables, urban heating incidents, leakage of pollutants, abrupt changes in vegetation, irrigation leakages, and weather extremes. All these phenomena would be recorded as fluctuations in a satellite signal. They may be expressed in different parts of the spectrum, evolve in different spatiotemporal scales, and they can influence the original signal without exceeding the range of normal values. They may occur regularly or unexpectedly, in known or unknown locations.

Usually, there is not enough information available on a fluctuation of interest to facilitate its isolation. In such cases, local fluctuations can be made more visible by suppressing patterns that are common to the majority of the dataset. Patterns may be suppressed by explicitly modelling and removing signal components, if characteristics of these components are *a priori* available. An alternative option would be the application of normalization techniques, which rescale the data, provide adjustment for overall patterns and allow the data to become internally comparable.

Popular choices to identify and remove general trends and periodic signal components include autoregressive and ordinary regression models and filtering and decomposition techniques (e.g. Cleveland et al., 1990; West, 1997; Jonsson and Eklundh, 2002; Grieser et al., 2002; Alegana et al., 2013; Wang et al., 2014). Wavelets and Fourier transforms are widely used to define signal patterns of different periodicity in a variety of geophysical applications, ranging from climatic studies to hazards and environmental research (e.g. Meyers et al., 1993; Kumar and Foufoula-Georgiou, 1997; Ghil et al., 2001; Sajda et al., 2002; de Jong and van der Meer, 2004; Labat, 2005; Scharlemann et al., 2008; Humlum et al., 2011; Pyayt et al., 2013; Tary et al., 2014; Qader et al., 2015). Randolph (2005) describes typical methods to normalize signals and/or images; the adjustments presented in his work include corrections for constant and non-constant shifts, scaling, and combinations thereof. Adaptations of the Standard Normal Variate method he mentions are often applied in geosciences to identify extremes (e.g. Tramutoli, 1998; Jimenez-Munoz et al., 2013, 2015).

These methods face three possible limitations. First of all, they cannot easily define signal components with a period longer than

* Corresponding author.
E-mail address: e.pavlidou@utwente.nl (E. Pavlidou).

http://dx.doi.org/10.1016/j.cageo.2016.02.016
0098-3004/© 2016 Elsevier Ltd. All rights reserved.

the temporal length of the available dataset. Secondly, large-scale patterns which are changing through time may still obscure fluctuations of very small intensity. A third limitation arises because existing methods rely only on the temporal dimension of the measurement. This is an issue especially relevant for research based on multidimensional data, as is the case of hyper-temporal satellite measurements. Satellite sensors provide synoptic, time-synchronous and consistent sampling of geophysical parameters over large areas and over long periods of time; both the temporal and spatial characteristics of the data are needed to extract fluctuations of these parameters from a complicated mix of different influences and noise. As noted by Tary et al. (2014), multi-dimensional geophysical data are traditionally analysed individually in a one-dimensional manner, but it would be very important to consider more dimensions in the analysis.

Spatial information is commonly used in satellite image processing. Such image processing techniques have been applied, among others, for feature extraction (e.g. van der Werff et al., 2006; Soto-Pinto et al., 2013), change detection (e.g. Coppin et al., 2004; Canty and Nielsen, 2012), normalization of specific influences in the satellite signal (e.g. Yang and Lo, 2000; Canty and Nielsen, 2008; Ulusoy et al., 2012), active fire monitoring (e.g. Giglio et al., 1999, 2003; Ichoku et al., 2003; Stolle et al., 2004; Kuenzer et al., 2007, 2008; Calle et al., 2008; Xu et al., 2010; Wooster et al., 2012), and studies of volcanic and geothermal activity (e.g. Coolbaugh et al., 2007; Ganci et al., 2011; Koeppen et al., 2011; Murphy et al., 2011; Steffke and Harris, 2011; Gutierrez et al., 2012; Vaughan et al., 2012; Blackett, 2014, 2015; van der Meer et al., 2014). A main aim is often the selection of an optimum background: this is used to contrast with an expected change or to describe a representative average state of a given neighbourhood. Subsequent processing and statistical analyses vary depending on the field of application.

Our work is a modification on aforementioned attempts to apply normalization for spatial data. We build on kernel-based approaches (for an overview see, e.g., Canty, 2010). We use a deconvolution matrix to select a pixel's spatial neighbourhood and normalize every pixel by the mean value of its neighbourhood set. In this way, patterns that are present in the central as well as the surrounding pixels are suppressed, and localized fluctuations are made more visible. We then apply a window-based retrieval technique to isolate these fluctuations in the normalized time series.

The added value of our approach is that spatial and spectral techniques are combined in a single algorithm and in a non-application-driven manner. Our normalization can suppress regionally extended patterns at different time-scales (diurnal, seasonal, yearly, etc.), based on both spatial and temporal components of the original data, and independently of the length of the dataset. Processing is run uniformly in the whole dataset and can isolate fluctuations which are not expected or known in advance. The processing chain is at the same time generic and flexible enough to be applied in different domains, and may be applied in near real-time mode.

We demonstrate our approach on a hyper-temporal geostationary thermal infra-red (TIR) dataset, recorded by the SEVIRI sensor onboard the Meteosat-9 satellite, and subset over two areas of different TIR variability in time and in space. We choose to base our case study on TIR data because of the wide range of earth-science related TIR applications: monitoring of fire and volcanic activity, geothermal exploration, and others (Sobrino et al., 2013; Ulusoy et al., 2012). The resolution of the sensor (3×3 km spatial, 15-min temporal) supports temporally intensive monitoring. In real-life applications, it rarely happens that the same well-known fluctuation is repeated in different conditions. Thus, to be able to evaluate the performance of our algorithm in different contexts and with better control, we carry out experiments with known synthetic fluctuations imposed in real data. We increase a small number of consecutive brightness temperature (BT) values in the original data and show that the normalization makes these increases more visible. We retrieve the fluctuations in the normalized time series and evaluate retrieval in reference to the intrinsic signal variability.

We then compare our findings with results of the application of widely used statistical and spectral methods. We decompose the TIR signal using Seasonal-Trend Decomposition based on LOESS (STL), following Cleveland et al. (1990) and Hafen (2010). In this way we remove the dominant daily and seasonal component of the data and we try to detect the artificially imposed fluctuation in the remainder. As an alternative, we use Fast Fourier Transform to define all principal frequencies of the signal. Similar work has been done, for example, by Humlum et al. (2011) to study periodic oscillations in climatic records and by Wang et al. (2014) to study water level fluctuations. We remove the defined patterns and reconstruct the signal based on the remaining frequency components, in order to trace the imposed fluctuation there. Finally, we apply the Standard Normal Variate method version which Jimenez-Munoz et al. (2013, 2015) used to detect anomalous thermal episodes over the Amazon. The technique rescales the data using their mean and standard deviation; thresholds are then applied to detect anomalous standardized values. The comparison between results of the different approaches concludes the performance evaluation of our method.

2. Methodology

In the first part (Section 2.1) we present our method. We then shortly present the three spectral and statistical approaches we applied to evaluate its performance (Section 2.2).

2.1. Proposed method

Our proposed method consists of three steps: pre-processing, normalization and retrieval (Fig. 1).

Pre-processing. Image pre-processing is intended to discard measurements that have been disturbed by factors other than the variable of interest. This step requires consistently defined *a priori* knowledge on the presence of disturbances. Pre-processing is specific to each application and dataset, and is not covered here in detail. In the case of most satellite imagery-based studies, the main factor interfering with the signal from the earth's surface is the atmosphere, and especially the presence of cloud cover.

Normalization. In this step we use spatial information to suppress temporal patterns common between a pixel and its neighbourhood, without explicitly modelling them. As a result, fluctuations which affect only the central pixel stand out.

We define a square, single-pixel-wide frame of neighbouring pixels (from here on, 'neighbourhood set, see Fig. 1A). The neighbourhood set lies at a defined distance around a central pixel (Fig. 1A, details on choosing the distance are provided later). The inside of the frame is not included in the neighbourhood set. The underlying assumption is that a localized fluctuation has a spatial extent smaller than the area framed by the neighbourhood set; it is thus contained in the central pixel, but not in its neighbours. We divide the central pixel value by the average value of the neighbourhood set (Fig. 1B). The process is repeated for all time-steps and results in a spatially based normalized time-series, which retains the temporal resolution and serial dependence of the original time-series.

The normalized time series shows the degree of dissimilarity between a pixel and its surroundings. In that sense, it expresses

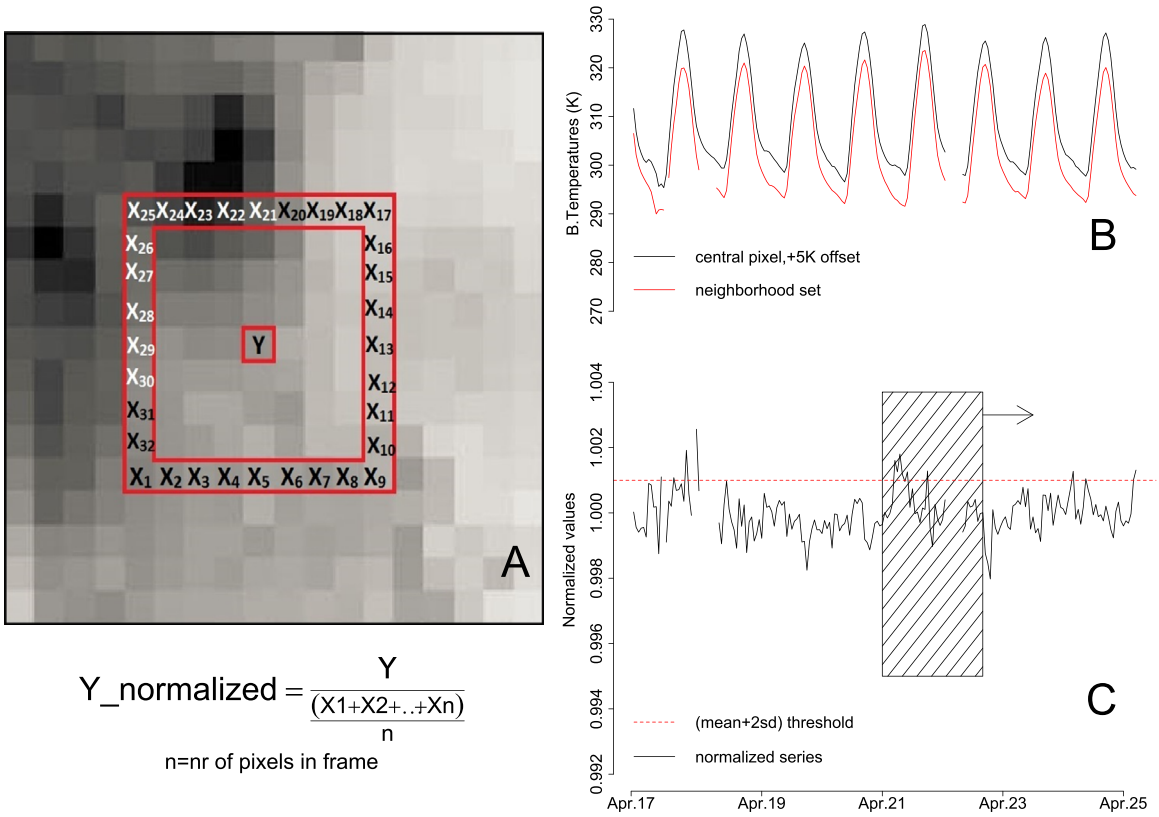

Fig. 1. (A) With normalization, every pixel value in the image is divided by the average value of a frame of neighbouring pixels. The pixels between the frame and the central pixel are not included to ensure that a localized fluctuation is not averaged out. (B) Original time series from central Pixel Y (black line) and the average value of its neighbourhood set (red line), with offset for clarity. (C) The same series after normalization. The daily pattern common in the central and neighbouring pixels is no longer dominant in the data. During retrieval, normalized values above the threshold are flagged. A temporal moving window counts the number of flags within a specific time period. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

local heterogeneity and the way this changes over time. When the central and neighbouring pixels are similar, the expected normalized value is approximately 1 and constant throughout time (Fig. 1C). Normalized values deviate from 1 as the differences between the central pixel and the frame increase. These differences reflect the natural variability in the signal, which may be due to differential heating of the earth surface, atmospheric effects, differences in reflective/emissive properties of different land covers, soil moisture, etc. An anomalous fluctuation in this context is defined as an unusual change in the relation between the pixel and its surroundings. This translates to normalized values largely deviating from the series mean.

Natural variability of the signal is statistically described by its standard deviation (σ). Normalization is designed to minimize this natural variability by suppressing signal patterns which are common between the pixel and its neighbours. Short-term localized anomalous fluctuations are better visible when the rest of the normalized series has minimal variability. We thus use σ of the normalized series to evaluate performance of normalization and to choose the distance between the normalization frame and the central pixel. We first calculate σ for each normalized time series separately. We then calculate the average σ of all normalized series in the dataset to describe variability for the whole study area ($\sigma_{dataset}$). We choose as optimal the normalization distance which results to the lowest $\sigma_{dataset}$. Once set, the distance between central pixel and frame remains the same for all the pixels of the dataset. This setting is not built in the code because every study area has different natural variability, potentially changing through time. By allowing the distance to be user-defined, we ensure flexibility of the method.

The statistical metric of $\sigma_{dataset}$ is the only criterion to choose frame size in the case of unknown fluctuations. However, it should be critically assessed as well, especially if there is more information available on an expected fluctuation. The distance should be large enough that a localized fluctuation is not included in the neighbourhood set. If the aim is to detect an oil spill extending over hundreds of km^2, it would be pointless to choose a 1-km distance (even if this resulted to the lowest $\sigma_{dataset}$). Larger sets increase statistical consistency in the normalization process and may decrease the variability. However, if the size of the set is so large that the central pixel and its neighbours are incomparable (e.g. if they belong to different climatic regions), there may be hardly any common patterns left for the normalization to suppress. In the cases mentioned above, the optimal choice would be the set which fulfills all provisions: lowest $\sigma_{dataset}$ possible; and, distance larger than the expected spatial extent of the fluctuation but within the limits of the same climatic region.

Data availability may influence the result of normalization. In each time-step different parts of the normalization set may be missing. In case of extensive cloud cover there is a possibility that the few pixels that remain after cloud-masking are not representative of the complete neighbourhood set. We set a threshold to ensure that normalized values of consequent time-steps are comparable. The threshold discards normalized values which have been calculated with less than a specific portion of the pixels of the neighbourhood set. This threshold is optional. We statistically test if its application is needed and define its level (see following sections).

Retrieval. Retrieval highlights and temporally isolates periods in time when the normalized signal is dominated by values strongly deviating from the mean. Fluctuations which are not only localized but also temporally persistent are retrieved using a window-based

approach. A temporal moving window can locate fluctuations in time and its application is independent of the process that may be generating the fluctuation (Chandola et al., 2012).

A *mean*+2σ threshold is applied on the normalized time series to distinguish the values which most strongly deviate from the mean. All values exceeding this threshold are flagged (Fig. 1C) and from here onwards they will be referred to as anomalies. Consequently, a temporal moving window counts the anomalies falling within the specified duration of the window (Fig. 1C). The temporal length of the moving window is defined by the user, depending on the application, the temporal resolution of the sensor and the desired level of detail in temporally locating the fluctuation. The resulting time series consists of the total number of anomalies per window and represents temporal clustering of highly deviating normalized values.

Also for this step the results of retrieval may be affected by data availability. Lack of data within the temporal window may result in low numbers of flags; these may be mistakenly considered to reflect low values in normalized data. To compensate for this effect, we divide the number of anomalies in the temporal window by the ratio *existing observations/theoretical number of observations*. Furthermore, we use a threshold to discard values which were calculated with less than a minimum number of available observations. This threshold is optional. We statistically test if its application is needed and define its level (see following sections for more details).

2.2. Traditional methods applied for evaluation

We apply three commonly used spectral and spatial approaches.

First we use Seasonal-Trend decomposition by LOESS (STL) (Cleveland et al., 1990; Hafen, 2010; Wang et al., 2014) to define the daily and yearly components in the data. These components were subtracted by the original values. We examine the remainder to detect the imposed +3 K fluctuation.

We then apply Fast Fourier Transform (FFT) to define all principal frequency components in the dataset (Humlum et al., 2011; Wang et al., 2014). FFT requires continuous series; to eliminate missing data, we apply Singular Spectrum Analysis (SSA). This is a gap-filling method which preserves periodic patterns of the signal (Buttlar et al., 2014; Korobeynikov, 2010; Kondrashov and Ghil, 2006). After performing forward-FFT, we remove the most dominant frequencies and use inverse-FFT to reconstruct the signal. We expect to detect the imposed +3 K increase in the reconstructed signal.

Finally, we apply the version of Standard Normal Variate method which was recently presented in Jimenez-Munoz et al. (2013, 2015). Following this approach, standard scores are calculated from original values to show how many standard deviations is the distance between each observation and the mean of the series:

$$BT_{standardized} = \frac{BT_{original} - BT_{mean}}{\sigma} \quad (1)$$

The authors classify standard scores to the following levels of warming (standardized score range in brackets): abnormal [+0.5, +0.8], moderate [+0.8,+1.3], severe [+1.3,+1.6], extreme [+1.6, +2.0] and exceptional [+2.0 and higher]. The probability of a score being anomalous is up to 57.6% for the abnormal warming level, up to 80.4% for moderate, up to 86% for severe, up to 95.4% for extreme and more than 95.4% for exceptional warming levels (Jimenez-Munoz et al., 2015).

We apply this method in the three recommended scales: monthly (June), seasonal (May, June and July), and yearly (2011).

3. Application and evaluation

In the first part of this section we demonstrate application of

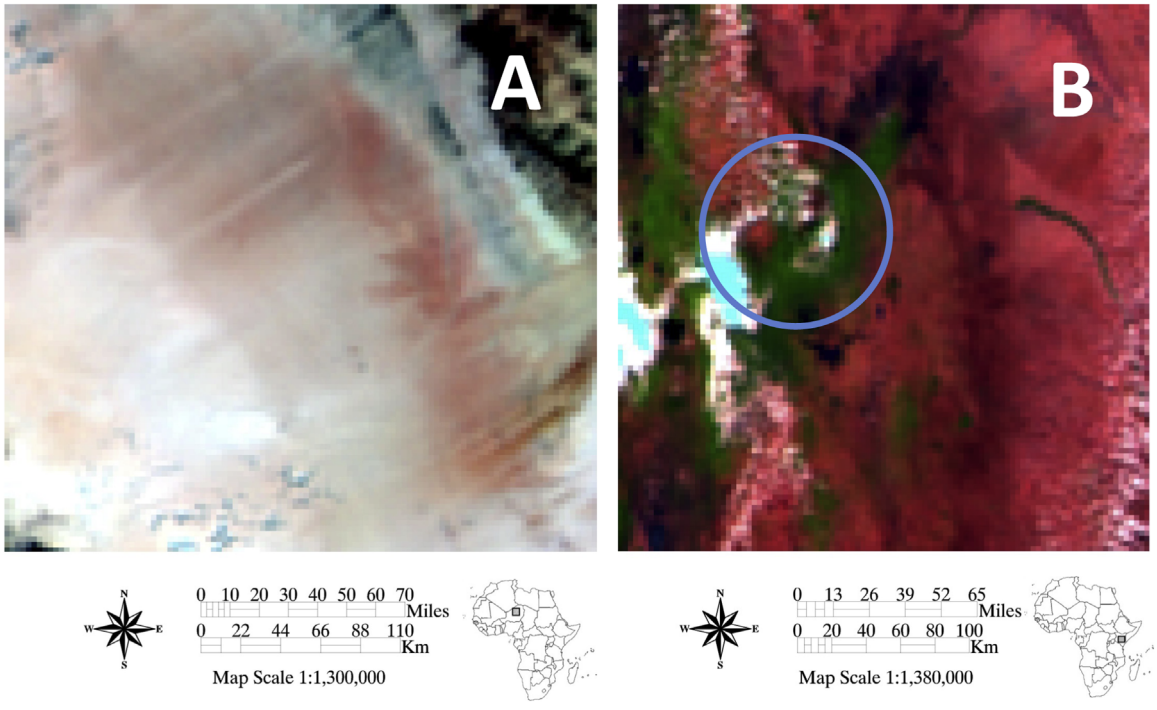

Fig. 2. Natural colour RGB images from the study areas in Niger (A) and Kenya (B). The area in Niger is almost completely covered by desert consisting of sand and gravel, with the exception of the rock formations in the NE corner of the image. Kenya has a more complex cover including bare soil, vegetated land, rock formations, water bodies, urban areas and mountains. The main feature is Mount Kenya circled in blue, in this image partly covered by clouds (white and turquoise). (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

Table 1
Experiments and related choices.

Objective	Experimental setting
Choice of settings	
1. Size of normalization set: Prefer the size which results to normalized series with the lowest standard deviation	Use frame side of 9, 17, 25, 33, 41, 49, 57, 65, 73 pixels, corresponding to approximate radii ranging between 13 and 108 km.
2. Retrieval with temporal windows of different length: Prefer windows of similar duration as the fluctuation of interest	5-day long artificial fluctuation retrieval with 2-, 4-, 7-, 10- and 20-day windows.
3. Data availability threshold for normalization: Choose the minimum data availability that preserves the distribution of values in the set	Normalize with sets of minimum 30%, 50%, 60%, 70%, 75%, 80%, and 100% data availability. Record standard deviation and distribution of values within the set.
4. Data availability threshold for retrieval: Choose the minimum data availability that preserves the distribution of values in the series leaving enough data for analysis	Retrieval with minimum 10%, 20%, 30%, 35%, 40%, 50%, 60%, and 80% of data available .
Experiments on normalization and retrieval	
5. Spatial extent of fluctuations suppressed by the normalization	Impose +3 K fluctuation in (a) 1 pixel, (b) 10 pixels and (c) 7411 pixels (67% of the image).
6. Evaluate intensity of retrieved fluctuations	Impose and retrieve fluctuations of +1, 2, 3, 4, 5, and 6 K.
7. Retrieve fluctuations in different locations in the image and in the time series	Impose and retrieve the same +3 K fluctuation. (a) In timeslots with/without gaps (February, June, September, and November) (b) At locations with high/low variability in the normalized series.
8. Effect of clouds on normalization	Process data with and without cloud masking.

our method on two spatial subsets of a year-long satellite TIR dataset. The two subsets represent study areas of different homogeneity (e.g. in local weather, land cover, and anthropogenic activities, see Fig. 2). We impose in both datasets artificial fluctuations in the form of increased brightness temperatures (see Table 1 for details). We describe how we defined the settings for the processing. We evaluate the choices in terms of their effect on single-series σ and $\sigma_{dataset}$. We then retrieve the imposed anomalies and evaluate performance of the method. In the second part of the section we process the same data with three traditional approaches and compare results with our method to evaluate it further.

3.1. Input

We used TIR imagery acquired from the Spinning Enhanced Visible and Infra-Red Imager (SEVIRI) onboard EUMETSAT's Meteosat-9 geostationary satellite, positioned at 0°/36,000 km. The instrument has a nominal spatial resolution of 3×3 km^2 at nadir and a sampling rate of 15 min in the TIR channels. We used two spatial subsets of a year-long (2011) whole-disk dataset from channel 9, registered at a wavelength range of 9.8–11.8 μm (λ_{cent} = 10.8 μm). Original top-of-atmosphere radiance values were converted to Brightness Temperatures (BT, in K, following Clerbaux, 2006). The first study area (spatial extent: 327×303 km or 109×101 pixels) was located in the desert in Niger, and served as an example of a homogeneous background with low spatial and temporal variability in the TIR. The second study area was located in Kenya (spatial extent 309×318 km or 103×106 pixels) and was very diverse in terms of geomorpology and land cover, representing a very heterogeneous background.

3.2. Pre-processing

We masked cloud-affected pixels using masks of EUMETSAT's Climate-Monitoring Satellite Application Facility (CM-SAF). This dataset was produced with software of the Nowcasting SAF (Derrien and Le Gleau, 2011). We excluded cloud-filled and cloud-contaminated pixels from further processing. As some of the clouds in the dataset were not detected by the available masks, we further discarded remaining pixels with values lower than the lowest recorded temperatures in historical archives.

3.3. Results and evaluation of our method

We first present and evaluate the choice of settings applied in the processing; then we evaluate performance of retrieval. All the performed tests are described in Table 1. Experiments 1–4 (Table 1) were designed to test the following settings: size of the neighbourhood set, data availability thresholds, and length of retrieval window. These settings are application-dependent; the choices we present are specific for this case study. Experiments 5–8 (Table 1) were designed to test performance of normalization and retrieval of synthetic fluctuations.

Settings. The first parameter to test was the optimal size of the normalization set (Experiment 1, Table 1). We performed normalization with different sizes of neighbourhood sets, their frame sides ranging from 9 to 73 pixels (respectively 12–108 km shortest distance between each side and the central pixel). We selected the size with the minimum $\sigma_{dataset}$. For Niger and Kenya the statistics were different but in both cases we had minimum $\sigma_{dataset}$ when applying a 57-pixel frame side. The $\sigma_{dataset}$ was increased from 0.001 (57-pixel frame side) to 0.003 (9-pixel frame side) in Niger, and from 0.003 to 0.006 respectively in Kenya (Fig. 3A).

We tested retrieval of a 5-day artificial fluctuation using temporal windows of 2–20 days. The fluctuation was retrieved as an anomaly in all of the cases. The main differences in retrieval were the timing, the width and the size of the resulting peak. For all windows the anomaly starts to be visible at the same time, but the center of the peak shifts to later times for larger windows. The peak is higher for longer windows since more anomalous values are captured in a single window. The optimum window length is roughly the same as the anticipated length of the anomaly. If this is not known in advance, it is recommended to test multiple windows and assess the persistence of peaks between short(er) and long(er) windows.

The application of data availability thresholds is not relevant for datasets which are not affected by missing values. We calculated the $\sigma_{dataset}$ at different data availability levels to see if the use of these thresholds would be needed in this study. We found that a 20% decrease in data availability doubled the $\sigma_{dataset}$ of Niger; in Kenya the corresponding $\sigma_{dataset}$ became five times higher. We thus decided that it would be relevant to apply the thresholds in order to minimize this variability. The choice on these thresholds is based on two criteria: (a) preserve the distribution of values in the normalization set/temporal window and (b) make sure that enough data remain for analysis. We decided that a minimum of 75% data was needed for normalization, and a corresponding 25% for retrieval. The reduction in variability of normalized series after application of a 75% threshold can be seen in Fig. 3B.

The examples we show in the figures are obtained with the following optimal settings: a normalization set of a 57-pixel frame

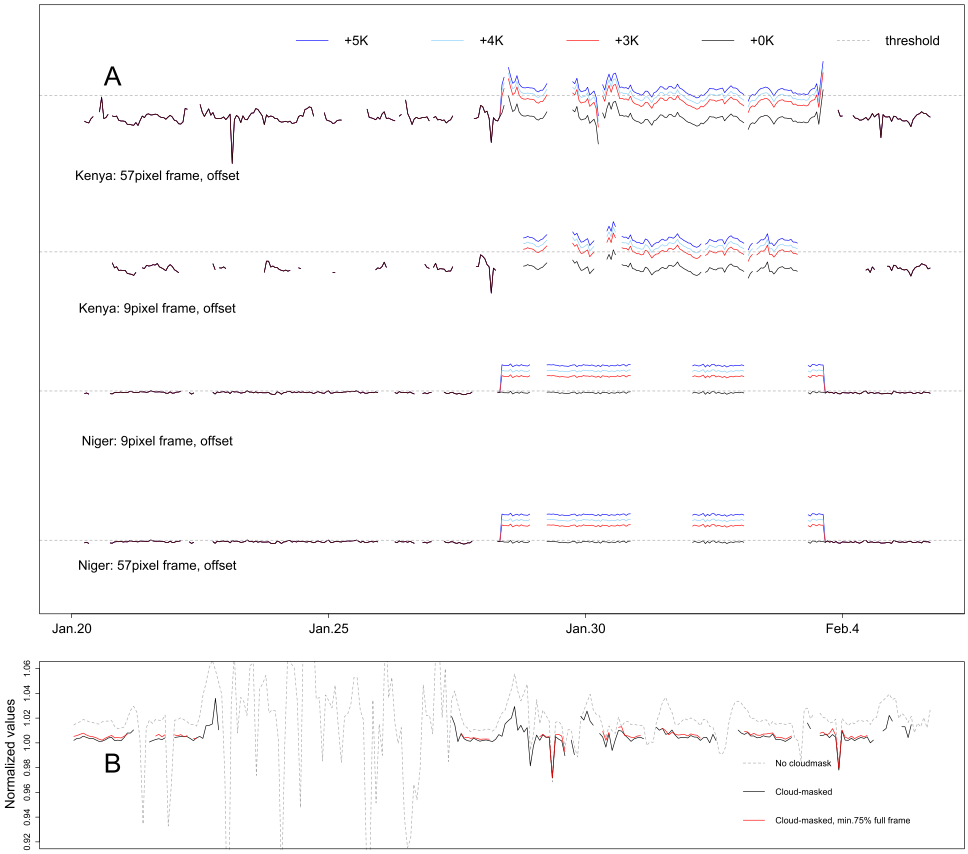

Fig. 3. Effect of algorithm settings and missing data. (A) Effect of frame size. The panel shows the result of normalization of series with different imposed fluctuations, using different frame sizes, in both study areas. The series are displayed with offset for clarity. Normalized series in the heterogeneous area of Kenya (first two rows) are more variable than in homogeneous Niger (last two rows). Imposed fluctuations of +(3–5) K exceed the detection threshold regardless of frame size, and are more visible when the variability is low. Artificial fluctuations increase linearly the normalized values. (B) Effect of clouds and missing data. The presence of clouds increases variability in the time series (grey dashed line); even after cloud masking, cloud remnants introduce spikes. Lack of data in the neighbourhood set also increases variability of cloud-masked normalized series. In contrast, the series normalized with at least 75% of the data available (red solid line) is the least variable. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

side, with at least 75% of the pixels present; and a 7-day long temporal window, with at least 25% of the observations present.

Retrieval performance. With Experiment 5 (Table 1) we tested the hypothesis that fluctuations of extent smaller than the normalization set are detectable. We imposed the same +3 K fluctuation in areas of different spatial extent. As expected, the fluctuation was visible when its extent was smaller than the area framed by the normalization set but not when its extent was larger. In Fig. 4, artificial fluctuations of +3 K which cover both the central pixel, coincide after retrieval with the +0 K imposed fluctuation (black line).

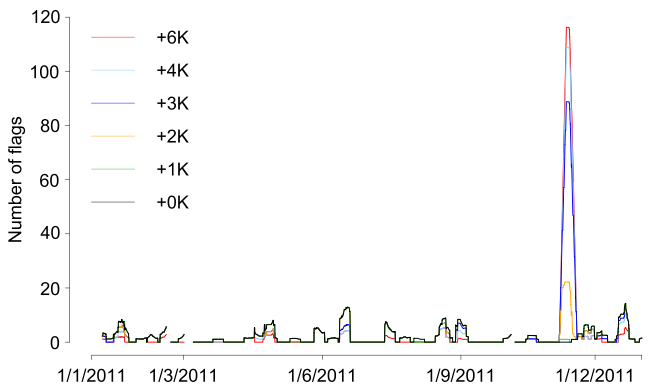

Fig. 4. Detection of synthetic anomalies of different magnitudes (Niger). Localized artificial fluctuations are retrieved as peaks of different magnitudes.

In Experiment 6 (Table 1) we imposed fluctuations of magnitude (+1) to (+6) K in the data (Fig. 4 for Niger). Regardless of the size of the fluctuation, higher fluctuations linearly increased normalized values (Fig. 3A). In the retrieval the effect was non-linear. The smallest imposed fluctuations we could retrieve, under conditions of high variability and low data availability, were +2 K in Niger and +3 K in Kenya (Fig. 4).

We imposed fluctuations in different locations in space and in time, to examine cases with higher and lower variability and with more or less missing values (Experiment 7, Table 1). Retrieved anomalies had different amplitudes when the same fluctuation was imposed in different times. Peaks were in general larger when the variability of the time series was lower and data availability was higher. As a result, we cannot quantify that a specific fluctuation will be always retrieved as an anomaly of corresponding amplitude; this should be evaluated in the context of each study area.

Our last experiment (Experiment 8, Table 1) tested the effect of cloud masking. Cloud masking reduced the variability of the normalized series approximately by a factor of 4 (Fig. 3B). The value of $\sigma_{dataset}$ changed from 0.020 to 0.005 after cloud masking in Niger and from 0.037 to 0.009 in Kenya. In contrast, the increase of $\sigma_{dataset}$ that was caused by localized fluctuations, like the ones we imposed in both datasets, was of the order of 10^{-4}. This shows the importance of removing disturbances from the data prior to processing.

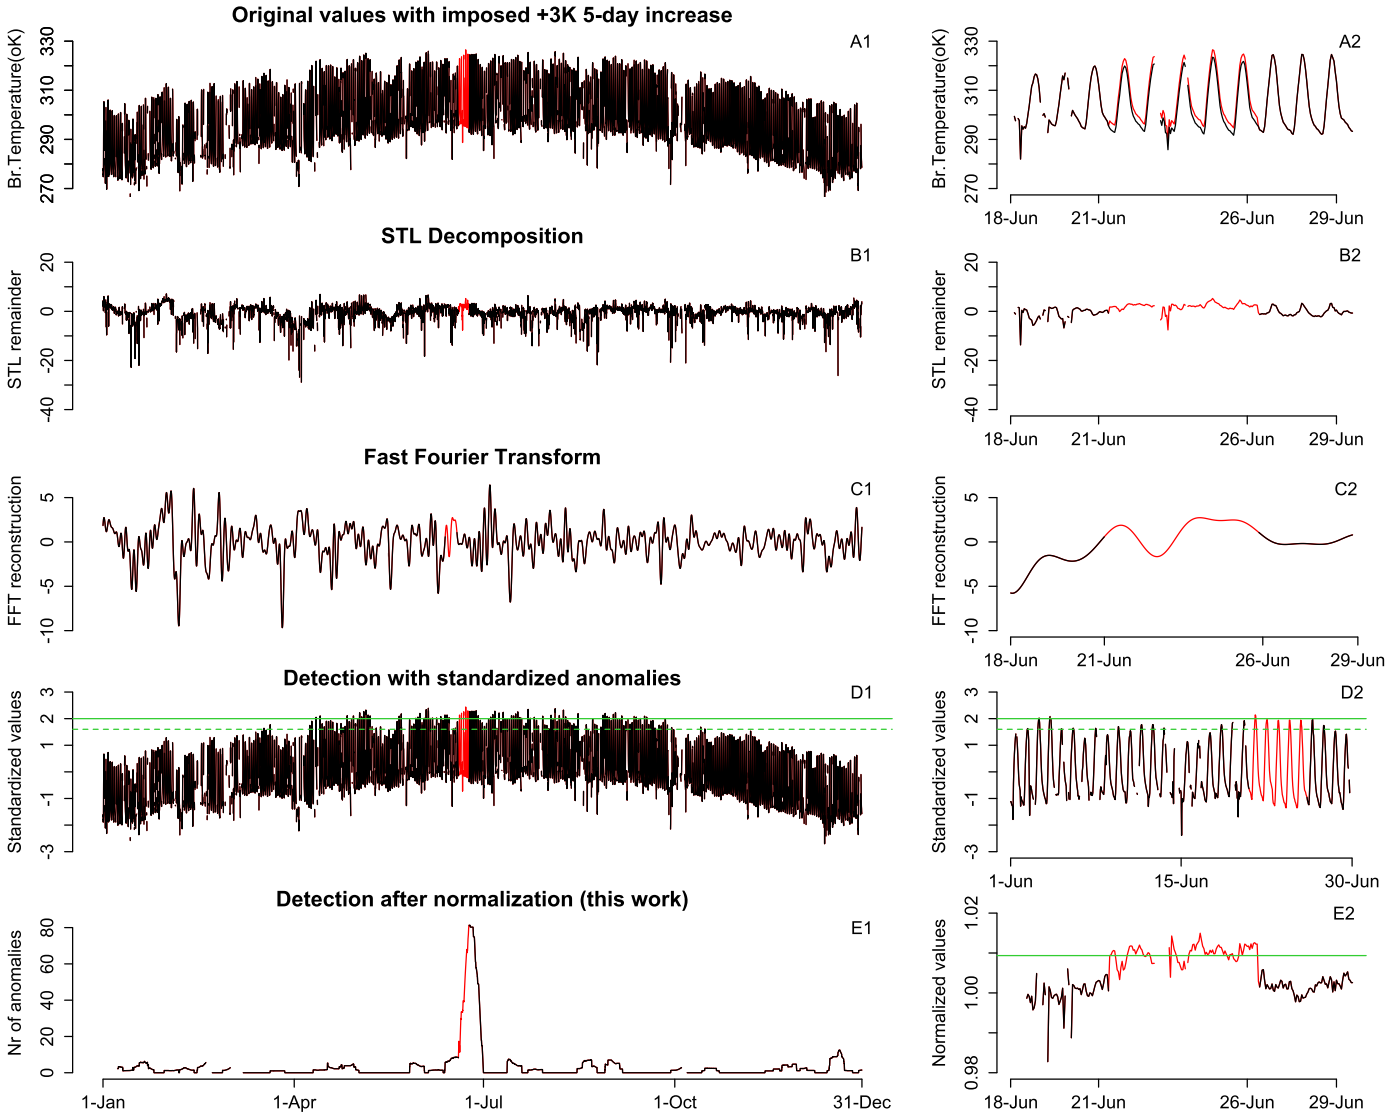

Fig. 5. Application of different methods to detect a synthetic +3 K 5-day BT increase imposed in the Niger dataset. A1: Original time series, with the imposed anomaly in red. The imposed anomaly does not stand out from the rest of the data. A2: the same (detail). B1: Remainder of STL decomposition of the original series. The anomaly is not distinguishable. B2: The same (detail). C1: FFT-reconstruction of the original series, after removing the three dominant frequency components (yearly, daily and twice-per-day). The impose anomaly does not stand out. C2: The same (detail). D1: Time series of the standard scores calculated from the original year-long dataset. The green dashed line designates extreme values, and the green solid line shows values which are classified as exceptional. The imposed BT increase is classified as anomalous but cannot be distinguished from the rest of the extreme/exceptional standardized values in the dataset. D2: The same but in monthly scale. E1: The methodology described in this paper clearly identifies the imposed anomaly as the single highest peak in the dataset. E2: Normalized dataset (detail). The imposed anomaly stands out in the normalized series, largely exceeding the detection threshold (in green). (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

3.4. Comparison with other methods

We used the original dataset from Niger and imposed a +3 K 5-day increase in BT in June 2011. This had proven straightforward to retrieve with our method, so we tested detecting it with others.

Results. Results from the application of all methods are summarized in Fig. 5.

The first row shows the original series with the imposed fluctuation.

The second row shows the remainder of STL decomposition. This was the result of subtracting the seasonal and daily patterns from the original series. The artificial fluctuation (in red) is contained in the remaining values, but cannot be distinguished from the rest of the data (Fig. 5B1,2).

In the case of FFT, the dominant frequency components were yearly, daily, and twice-per-day. We removed these components and reconstructed the signal using inverse FFT. The imposed anomaly did not stand out in the reconstructed signal (Fig. 5C1,2).

Detection with standardized scores in a yearly scale yielded extreme and exceptional anomalies throughout the period between mid-April to October (Fig. 5D1). Standardized data retained the seasonal pattern of the original dataset. The imposed fluctuation was included in the values which are considered extreme/exceptional, but so were many other values. As a result, the imposed anomaly did not stand out. We obtained similar results with analysis on seasonal and monthly scale (Fig. 5D2).

In contrast, after applying the normalization we describe in this paper, we detected the imposed +3 K 5-day increase as the single most prominent peak in the whole dataset (Fig. 5E1). The majority of anomalous values exceeded the detection threshold already in the normalized series (Fig. 5E2).

4. Discussion

Fluctuations of geophysical parameters are of interest in applications ranging from climatic studies to hazard monitoring. Detection of such fluctuations is often a complicated task, even with the variety of available satellite sensors. Depending on the parameter of interest, input from multiple parts of the spectrum may be needed; issues of saturation may arise in different scales; and fluctuations may be obscured by predominant patterns.

Localized fluctuations of interest may remain within the expected range of original values. In such cases, pixel values are normal in absolute value, and the fluctuation is easily obscured by predominant (daily and seasonal) patterns of the signal. Such fluctuations can be detectable if we examine each value in its spatial context (Byun et al., 2007). For example, a forest fire which would saturate the TIR channels of high-resolution sensors may increase the pixel value of low-resolution TIR imagery only by a few degrees Kelvin (Wooster et al., 2005). This increase is localized and alters the usual relationship between the affected pixel and its surroundings. Our methodology targets such fluctuations by identifying anomalies as deviations from the usual relation between a pixel and its spatial neighbourhood. This approach facilitates detection of contextual anomalies which remain within the normal range of values. Furthermore, it allows detection of fluctuations over constantly elevated or constantly low background. This could be the case, for example, of temporally variable volcanic activity in a permanent lava lake.

Considering the spatial context of an anomaly offers another advantage: it constrains its occurrence locally, providing insight on the relation between the anomaly and the potential underlying causes. An anomaly that extends in thousands of km^2 can most likely not be attributed to a spatially limited process. In that respect, the shape of the normalization frame serves as an upper boundary for the areal extent of a detected anomaly. The use of an open frame ensures that only a spatially limited fluctuation is detectable; if the fluctuation was included in the normalization set, it would be averaged out when we divide an anomalous pixel by equally anomalous surroundings. As we show in Experiment 5 (Section 3.3), anomalies are only detected when they are not covered by the frame.

In terms of the temporal dimension of the data, our method suppresses patterns not by modelling past observations, but by identifying the commonalities between signals of neighbouring areas in the time when they emerge. This is especially important because predominant patterns, like seasonality, often vary in time. For example, weather extremes are becoming more frequent (Easterling et al., 2000); by using our method, a short-lived localized brightness temperature increase caused by spatially finite surface processes can be distinguished from a short-term, unusually warm period affecting the whole TIR image. Present-time normalization of commonalities allowed our method to outperform approaches which do not fully address predominant patterns or rely on past observations to suppress them. Another advantage is that processing can take place in near-real-time mode. Additionally, by defining the temporal occurrence of an anomaly we can better evaluate its relation with potential causative processes.

Rather than pre-defining the settings of our method, we statistically determine specific criteria for the choices. In this way we ensure flexibility of the method and optimal application in study areas with different local conditions. Information on the characteristics of the fluctuation of interest can be used to constrain application of the methodology, but it is not required to have *a priori* knowledge of the expected fluctuation or the predominant patterns in the signal. Settings are applied uniformly to the whole dataset, and retrieval is automated, supporting objectivity and comparability of the results. The $mean+2\sigma$ threshold that isolates anomalies is based on each single-pixel time series; that has the advantage that anomaly detection is performed considering the local conditions. Normalized series of high variability has higher $mean+2\sigma$ threshold; then, values need to be higher before they are declared anomalous. Datasets of long duration are recommended to better describe the usual relation between a pixel and its surroundings.

Our results show that we could detect short-term (lasting 1.4% of the duration of the dataset), localized (affecting only one pixel) and low-intensity anomalies (as low as 0.7% of the signal, which was the case of +2 K over a 300 K background in Niger). This facilitates use of low spatial resolution sensors for detection of small-scale environmental changes, even with only one band available. The case study we presented as an example was based on geostationary TIR input. Geostationary products offer the advantage of high temporal coverage, which is important for timely monitoring of extremes (Giglio et al., 2003). We emphasize the use of hyper-temporal datasets further because temporal resolution is required to detect short-lived fluctuations. Our method, however, is not restricted to the input of our example; it may be applied at different scales and different types of spatio-temporal data. Higher spatial resolution may also provide more detail in locating a potential source of anomaly.

5. Conclusions

We focus on unexpected, localized, short-term signal fluctuations and propose a methodology to detect them using single-band input. Our normalization procedure suppresses spatially extended, large-scale temporal patterns in single-pixel time series without having to explicitly model them. Data are brought to scale and localized fluctuations with an extent smaller than the defined become more obvious, regardless of the underlying causative processes. The fluctuations are distinguished from large-scale periodical patterns by analysing both the spatial and temporal dimensions of geophysical data. This can serve a wide spectrum of applications and facilitate monitoring of extremes.

Acknowledgements

We would like to thank EUMETSAT and the Climate Monitoring SAF for providing datasets. We would also like to thank Petra E. Budde, Giovanni Buzzo, Wim H. Bakker, Nadira Khan, and Marisol Lopez-Meda for providing technical support and for their input in the development of the method. Finally we want to thank the reviewers, whose comments helped us to considerably improve our initial manuscript.

References

Alegana, V., Atkinson, P., Wright, J., Kamwi, R., Uusiku, P., Katokele, S., Snow, R., Noor, A., 2013. Estimation of malaria incidence in northern Namibia in 2009 using Bayesian conditional-autoregressive spatial temporal models. Spat. Spatio-Temporal Epidemiol. 7, 25–36. http://dx.doi.org/10.1016/j.sste.2013.09.001.

Blackett, M., 2014. Early analysis of Landsat-8 Thermal Infrared sensor imagery of volcanic activity. Remote Sens. 6 http://dx.doi.org/10.3390/rs6032282.

Blackett, M., 2015. An initial comparison of the thermal anomaly detection products of MODIS and VIIRS in their observation of Indonesian volcanic activity. Remote Sens. Environ. 171, 75–82. http://dx.doi.org/10.1016/j.rse.2015.10.002.

Buttlar, J., Zscheischler, J., Mahecha, M., 2014. An extended approach for spatio-temporal gapfilling: dealing with large and systematic gaps in geoscientific datasets. Nonlinear Process. Geophys. 21 http://dx.doi.org/10.5194/npg-21-203-2014.

Byun, Y.G., Huh, Y., Yu, K., Kim, Y.I., 2007. Evaluation of graph-based analysis for forest fire detections. Int. J. Comput. Electr. Autom. Control Inf. Eng. 1.

Calle, A., Gonzalez Alonso, F., Merino de Miguel, S., 2008. Validation of active forest

fires detected by MSG-SEVIRI by means of MODIS hot spots and AWiFS images. Int. J. Remote Sens. 29, 3407–3415. http://dx.doi.org/10.1080/01431160701596164.
Canty, M., 2010. Image Analysis, Classification, and Change Detection in Remote Sensing: With Algorithms for ENVI/IDL. Taylor and Francis, Boca Raton.
Canty, M., Nielsen, A., 2008. Automatic radiometric normalization of multitemporal satellite imagery with the iteratively re-weighted MAD transformation. Remote Sens. Environ. 112, 1025–1036. http://dx.doi.org/10.1016/j.rse.2007.07.013.
Canty, M.J., Nielsen, A.A., 2012. Linear and kernel methods for multivariate change detection. Comput. Geosci. 38, 107–114. http://dx.doi.org/10.1016/j.cageo.2011.05.012.
Chandola, V., Banerjee, A., Kumar, V., 2012. Anomaly detection for discrete sequences: a survey. IEEE Trans. Knowl. Data Eng. 24, 823–839. http://dx.doi.org/10.1109/TKDE.2010.235.
Clerbaux, N., 2006. Imager Processing: Data Dictionary. Technical Report MSG-RMIB-GE-TN-0024 version 2.2. Royal Meteorological Institute of Belgium.
Cleveland, R., Cleveland, W., McRae, J., Terpenning, I., 1990. STL: a seasonal-trend decomposition procedure based on LOESS. J. Off. Stat. 6, 3–73.
Coolbaugh, M., Kratt, C., Fallacaro, A., Calvin, W., Taranik, J., 2007. Detection of geothermal anomalies using advanced spaceborne thermal emission and reflection radiometer (ASTER) thermal infrared images at Bradys Hot Springs, Nevada, USA. Remote Sens. Environ. 106, 350–359. http://dx.doi.org/10.1016/j.rse.2006.09.001.
Coppin, P., Jonckheere, I., Nackaerts, K., Muys, B., Lambin, E., 2004. Digital change detection methods in ecosystem monitoring: a review. Int. J. Remote Sens. 25, 1565–1596. http://dx.doi.org/10.1080/0143116031000101675.
Derrien, M., Le Gleau, H., 2011. MSG/SEVIRI cloud mask and type from SAFNWC. Int. J. Remote Sens. 26 (21), 4707–4732. http://dx.doi.org/10.1080/01431160500166128.
Easterling, D., Meehl, G., Parmesan, C., Changnon, S., Karl, T., Mearns, L., 2000. Climate extremes: observations, modeling, and impacts. Science 289, 2068–2074. http://dx.doi.org/10.1126/science.289.5487.2068.
Ganci, G., Vicari, A., Fortuna, L., Del Negro, C., 2011. The HOTSAT volcano monitoring system based on combined use of SEVIRI and MODIS multispectral data. Ann. Geophys. 54 http://dx.doi.org/10.4401/ag-5338.
Ghil, M., Allen, M., Dettinger, M., Ide, K., Kondrashov, D., Mann, M., Robertson, A., Saunders, A., Tian, Y., Varadi, F., Yiou, P., 2001. Advanced spectral methods for climatic time series. Rev. Geophys. 11, 327–338.
Giglio, L., Descloitres, J., Justice, C., Kaufman, Y., 2003. An enhanced contextual fire detection algorithm for MODIS. Remote Sens. Environ. 87, 273–282. http://dx.doi.org/10.1016/S0034-4257(03)00184-6.
Giglio, L., Kendall, J., Justice, C., 1999. Evaluation of global fire detection algorithms using simulated AVHRR infrared data. Int. J. Remote Sens. 20, 1947–1985. http://dx.doi.org/10.1080/014311699212290.
Grieser, J., Tromel, S., Schonwiese, C., 2002. Statistical time series decomposition into significant components and application to European temperature. Theor. Appl. Climatol. 71, 171–183. http://dx.doi.org/10.1007/s007040200003.
Gutierrez, F., Lemus, M., Parada, M., Benavente, O., Aguilera, F., 2012. Contribution of ground surface altitude difference to thermal anomaly detection using satellite images: application to volcanic/geothermal complexes in the Andes of central Chile. J. Volcanol. Geotherm. Res. 237–238, 69–80. http://dx.doi.org/10.1016/j.jvolgeores.2012.05.016.
Hafen, R., 2010. Local Regression Models: Advancements, Applications, and New Methods (Ph.D. thesis). Purdue University.
Humlum, O., Solheim, J.-E., Stordahl, K., 2011. Spectral analysis of the Svalbard Temperature Record 1912–2010. Adv. Meteorol. 2011, Article ID 175296, http://dx.doi.org/10.1155/2011/175296.
Ichoku, C., Kaufman, Y.J., Giglio, L., Li, Z., Fraser, R.H., Jin, J.-Z., Park, W.M., 2003. Comparative analysis of daytime fire detection algorithms using AVHRR data for the 1995 fire season in Canada: perspective for MODIS. Int. J. Remote Sens. 24, 1669–1690. http://dx.doi.org/10.1080/01431160210144697.
Jimenez-Munoz, J., Mattar, C., Sobrino, J., Malhi, Y., 2015. A database for the monitoring of thermal anomalies over the Amazon forest and adjacent intertropical oceans. Sci. Data 2 http://dx.doi.org/10.1038/sdata.2015.24.
Jimenez-Munoz, J., Sobrino, J., Mattar, C., Malhi, Y., 2013. Spatial and temporal patterns of the recent warming of the Amazon forest. J. Geophys. Res.: Atmos. 118, 5204–5215. http://dx.doi.org/10.1002/jgrd.50456.
de Jong, S., van der Meer, F. (Eds.), 2004. Remote Sensing Image Analysis: Including the Spatial Domain. Springer, Netherlands.
Jonsson, P., Eklundh, L., 2002. Seasonality extraction by function fitting to time-series of satellite sensor data. IEEE Trans. Geosci. Remote Sens. 40, 1824–1832. http://dx.doi.org/10.1109/TGRS.2002.802519.
Koeppen, W., Pilger, E., Wright, R., 2011. Time series analysis of infrared satellite data for detecting thermal anomalies: a hybrid approach. Bull. Volcanol. 73, 577–593. http://dx.doi.org/10.1007/s00445-010-0427-y.
Kondrashov, D., Ghil, M., 2006. Spatio-temporal filling of missing points in geophysical data sets. Nonlinear Process. Geophys. 23.
Korobeynikov, A., 2010. Computation- and space-efficient implementation of SSA. Stat. Interface 3.
Kuenzer, C., Hecker, C., Zhang, J., Wessling, S., Wagner, W., 2008. The potential of multidiurnal MODIS thermal band data for coal fire detection. Int. J. Remote Sens. 29, 923–944. http://dx.doi.org/10.1080/01431160701352147.
Kuenzer, C., Zhang, J., Li, J., Voigt, S., Mehl, H., Wagner, W., 2007. Detecting unknown coal fires: synergy of automated coal fire risk area delineation and improved thermal anomaly extraction. Int. J. Remote Sens. 28, 4561–4585. http://dx.doi.org/10.1080/01431160701250432.
Kumar, P., Foufoula-Georgiou, E., 1997. Wavelet analysis for geophysical applications. Rev. Geophys. 35, 385–412. http://dx.doi.org/10.1029/97RG00427.
Labat, D., 2005. Recent advances in wavelet analyses: Part 1. A review of concepts. J. Hydrol. 314, 275–288. http://dx.doi.org/10.1016/j.jhydrol.2005.04.003.
van der Meer, F., Hecker, C., van Ruitenbeek, F., van der Werff, H., de Wijkerslooth, C., Wechsler, C., 2014. Geologic remote sensing for geothermal exploration: a review. Int. J. Appl. Earth Obs. Geoinf. 33, 255–269. http://dx.doi.org/10.1016/j.jag.2014.05.007.
Meyers, S., Kelly, B., O'Brien, J., 1993. An introduction to wavelet analysis in oceanography and meteorology: with application to the dispersion of Yanai waves. Mon. Weather Rev. 121, 2858–2866, doi: http://dx.doi.org/10.1175/1520-0493(1993)1212858:AITWAI2.0.CO;2http://dx.doi.org/10.1175/1520-0493(1993)121<2858:AITWAI>2.0.CO;2.
Sajda, P., Laine, A., Zeevi, Y., 2002. Multi-resolution and wavelet representations for identifying signatures of disease. Dis. Markers 18, 339–363. http://dx.doi.org/10.1155/2002/108741.
Scharlemann, J., Benz, D., Hay, S., Purse, B., Tatem, A., Wint, W., Rogers, D., 2008. Global data for ecology and epidemiology: a novel algorithm for temporal Fourier processing MODIS data. PLoS One 3 http://dx.doi.org/10.1371/journal.pone.0001408.
Sobrino, J., del Frate, F., Drusch, M., Jimenez-Munoz, J., Manuta, P., 2013. Review of high resolution thermal infrared applications and requirements: the Fuegosat synthesis study. In: Kuenzer, C., Dech, S. (Eds.), Thermal Infrared Remote Sensing: Sensors, Methods, Applications. Springer, Netherlands, pp. 197–214. http://dx.doi.org/10.1007/978-94-007-6639-6_10.
Soto-Pinto, C., Arellano-Baeza, A., Sanchez, G., 2013. A new code for automatic detection and analysis of the lineament patterns for geophysical and geological purposes (ADALGEO). Comput. Geosci. 57, 93–103. http://dx.doi.org/10.1016/j.cageo.2013.03.019.
Steffke, A., Harris, A., 2011. A review of algorithms for detecting volcanic hot spots in satellite infrared data. Bull. Volcanol. 73, 1109–1137. http://dx.doi.org/10.1007/s00445-011-0487-7.
Stolle, F., Dennis, R.A., Kurniwan, I., Lambin, E.F., 2004. Evaluation of remote sensing-based active fire datasets in indonesia. Int. J. Remote Sens. 25, 471–479. http://dx.doi.org/10.1080/01431160310001618022.
Tary, J., Herrera, R., Han, J., van der Baan, M., 2014. Spectral estimation—What is new? What is next? Rev. Geophysics 52, 723–749. http://dx.doi.org/10.1002/2014RG000461.
Tramutoli, V., 1998. Robust AVHRR techniques (RAT) for environmental monitoring: theory and applications. In: Cecchi, G., Zilioli, E. (Eds.), Proceedings of SPIE 3496. Earth Surface Remote Sensing II vol. 3496; 1998, pp. 101–113. http://dx.doi.org/10.1117/12.332714.
Ulusoy, I., Labazuy, P., Aydar, E., 2012. Stcorr: an IDL code for image based normalization of lapse rate and illumination effects on nighttime TIR imagery. Comput. Geosci. 43, 63–72. http://dx.doi.org/10.1016/j.cageo.2012.02.012.
Vaughan, G., Keszthelyi, L., Lowenstern, B., Jaworowski, C., Heasler, H., 2012. Use of ASTER and MODIS thermal infrared data to quantify heat flow and hydrothermal change at Yellowstone National Park. J. Volcanol. Geotherm. Res. 233–234, 72–89. http://dx.doi.org/10.1016/j.jvolgeores.2012.04.022.
Wang, F., Wang, X., Zhao, Y., Yang, Z., 2014. Long-term periodic structure and seasonal-trend decomposition of water level in Lake Baiyangdian, Northern China. Int. J. Environ. Sci. Technol. 11, 327–338. http://dx.doi.org/10.1007/s13762-013-0362-5.
van der Werff, H., Bakker, W., van der Meer, F., Siderius, W., 2006. Combining spectral signals and spatial patterns using multiple Hough transforms: an application for detection of natural gas seepages. Comput. Geosci. 32, 1334–1343. http://dx.doi.org/10.1016/j.cageo.2005.12.003.
West, M., 1997. Time series decomposition. Biometrika 84, 489–494. http://dx.doi.org/10.1093/biomet/84.2.489.
Wooster, M., Roberts, G., Perry, G., Kaufman, Y., 2005. Retrieval of biomass combustion rates and totals from fire radiative power observations: FRP derivation and calibration relationships between biomass consumption and fire radiative energy release. J. Geophys. Res. 110 http://dx.doi.org/10.1029/2005JD006318.
Wooster, M., Xu, W., Nightingale, T., 2012. Sentinel-3 SLSTR active fire detection and FRP product: pre-launch algorithm development and performance evaluation using MODIS and ASTER datasets. Remote Sens. Environ. 120, 236–254. http://dx.doi.org/10.1016/j.rse.2011.09.033.
Xu, W., Wooster, M., Roberts, G., Freeborn, P., 2010. New GOES imager algorithms for cloud and active fire detection and fire radiative power assessment across North, South and Central America. Remote Sens. Environ. 114, 1876–1895. http://dx.doi.org/10.1016/j.rse.2010.03.012.
Yang, X., Lo, C., 2000. Relative radiometric normalization performance for change detection from multi-date satellite images. Photogram. Eng. Remote Sens. 66, 967–980.