

Anomaly Attribution of Multivariate Time Series using Counterfactual Reasoning

Violeta Teodora Trifunov
Computer Vision Group
Friedrich Schiller University Jena
Jena, Germany
violetateodora.trifunov@uni-jena.de

Maha Shadaydeh
Computer Vision Group
Friedrich Schiller University Jena
Jena, Germany
maha.shadaydeh@uni-jena.de

Björn Barz
Computer Vision Group
Friedrich Schiller University Jena
Jena, Germany
bjoern.barz@uni-jena.de

Joachim Denzler
Computer Vision Group
Friedrich Schiller University Jena
Michael Stifel Center Jena for Data-Driven and Simulation Science
German Aerospace Center (DLR), Institute for Data Science
Jena, Germany

Abstract—There are numerous methods for detecting anomalies in time series, but that is only the first step to understanding them. We strive to exceed this by explaining those anomalies. Thus we develop a novel attribution scheme for multivariate time series relying on counterfactual reasoning. We aim to answer the counterfactual question of would the anomalous event have occurred if the subset of the involved variables had been more similarly distributed to the data outside of the anomalous interval. Specifically, we detect anomalous intervals using the Maximally Divergent Interval (MDI) algorithm, replace a subset of variables with their in-distribution values within the detected interval and observe if the interval has become less anomalous, by re-scoring it with MDI. We evaluate our method on multivariate temporal and spatio-temporal data and confirm the accuracy of our anomaly attribution of multiple well-understood extreme climate events such as heatwaves and hurricanes.

Index Terms—anomaly attribution, multivariate time series, counterfactual reasoning

I. INTRODUCTION

Finding causes of extreme weather events, power outages and abnormal fluctuations in financial data can be of crucial importance for their understanding and taking precautionary measures or even preventing them from occurring again in the future. We propose a novel anomaly attribution scheme to analyze anomalous intervals of multivariate temporal and spatio-temporal data and attribute those anomalies to a set of involved variables. To achieve this, we engage in answering the counterfactual question: would the anomalous event have occurred if one or more of the involved variables had been more similarly distributed to the data outside of the anomalous interval. For example, would there have been a heatwave in Europe, had the air temperature been lower in the summer that year or more similar to the last year’s summer’s average. However, since this question cannot be answered directly, we do the next best thing and replace the subset of variables in the detected anomalous interval by a random sample from their nominal distribution, as illustrated in Fig. 1 for a single

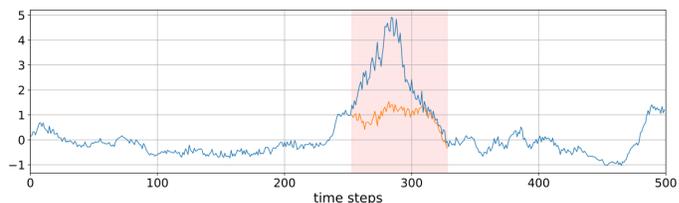


Fig. 1: Counterfactual replacement of the wave height (H_s) variable (orange) within the anomalous interval highlighted in red and the data before the replacement (blue).

variable. Our novel replacement algorithm preserves both the inter-variable and the inter-temporal correlations of the data.

We then examine whether the same interval became less anomalous after replacing a specific subset of variables by applying the Maximally Divergent Interval (MDI) method [1] for anomaly detection, repeatedly on each of the subsets. By determining which variables yield the lowest anomaly score after the replacement, we can conclude that those variables were responsible for the anomaly. We are only interested in attributing the anomaly to the subset of variables of the time series we are analysing and do not consider what caused those variables to be anomalous. However, in all our experiments, we examine a certain time frame before the detected anomaly in order to check if any other variable of the time series behaved abnormally and possibly influenced the anomalous event.

Our attribution method can be applied to any multivariate time series data regardless of potential outliers and missing values. This also makes it appropriate for the task of data repairing. Furthermore, it is one of the few attribution schemes for multivariate time series, and offers interpretation of the MDI’s anomaly detections.

The remainder of the paper is organized as follows. We discuss other time series anomaly attribution approaches in Section II. In Section III, we introduce the basic concepts

of causality and counterfactual reasoning and briefly review the MDI algorithm for anomaly detection. Our novel anomaly attribution approach for multivariate time series is introduced in Section IV. We evaluate our method by attributing anomalies on multiple datasets of ecological time series and the spatio-temporal climate data where the causes of the anomalies are well-understood. The results are presented in Section V and Section VI concludes this paper.

II. RELATED WORK

With the increase of the available data, creating efficient methods which can analyze it, detecting anomalies, and explaining what caused the anomalies as reliably as possible becomes crucial. There are numerous anomaly detection approaches for time series [2], but not many are concerned with attributing those anomalies. One of them, however, is the work by Siddiqui et al. [3], where the authors introduce a time series anomaly attribution method called TSInsight, which uses an auto-encoder with a sparsity-inducing norm on its output to the classifier. It learns to keep the features that are important for prediction by the classifier and discards the unimportant ones. TSInsight can create both instance-based and model-based explanations. This feature attribution method, unlike our approach, does not use counterfactual variables to attribute anomalies. Moreover, it does not attribute the anomalies to variables in the input data directly but to features derived from them. However, similar to our idea, Idé et al. [4] propose to explain the anomalous prediction produced by a black-box regression model by inferring the responsibility score of each of the input variables. They introduce the Likelihood Compensation (LC) method, based on the likelihood principle, i.e., on the proposition that, given a statistical model, all the evidence in a sample that is pertinent to the model parameters is enclosed in the likelihood function. This method computes a correction to each input variable. In contrast to our approach, however, the LC method does not consider replacing certain intervals of a subset of variables and is only applied to anomalous forecasts. Using MDI, we can detect any anomalous interval in the entire time series and correct and attribute it using our method, regardless of any outliers or missing values. Next, Zhang et al. [5] also engage in iterative repairing of anomalous time series data for improving pattern mining or classification tasks. Unlike our method, it does not aim to provide answers to what caused the anomalies themselves and is more computationally intensive. In the work of Shadaydeh et al. [6], anomalies are attributed based on the changes in the spectral cause-effect relationships between each pair of the involved variables. Unlike our method, it does not identify the subset of variables that contribute most to the anomalous event. Furthermore, in the method for attribution of multivariate extreme events proposed by Guaniche et al. [7], the authors aim to answer the question of how much does each variable contribute to the Mahalanobis distance metric. However, this method is applied on each time step independently rather than interval-wise as is the case with our method.

III. METHODOLOGICAL BACKGROUND

We will now briefly discuss the fundamental concepts of causality and counterfactual inference, followed by a more in-depth explanation of the method we use for detecting the anomalous intervals of multivariate time series to which our novel attribution scheme is tailored.

A. Causal and Counterfactual Reasoning

To better understand the world, humans seek to uncover the causes of certain events or phenomena by observing them over time. However, for different reasons, the answer to the question of why something happened is sometimes difficult to obtain. One way to empirically determine a cause of a particular event is to alter one of its potential causes at a time and observe if the alteration changed the event in question. This alteration is formally known as *intervention*. However, when the outcome of an event of interest cannot be reproduced due to the intervention being impractical or unethical, one needs to turn to the *counterfactual reasoning*. Namely, one can perform a thought experiment, asking oneself how the world would have changed had certain different actions been taken. In our case, it is impractical to intervene on climate variables and observe what kind of intervention would make a specific event less anomalous. Still, we can answer the question of how the anomalous event would have changed had a specific subset of the involved variables been distributed like the rest of the data outside of the anomaly. When this type of intervention results in a less anomalous event, we attribute it to the subset of the variables on which we intervened.

B. MDI Algorithm for Anomaly Detection

The MDI [1] is an unsupervised method for detecting anomalies in multivariate temporal and spatio-temporal data. Following the notation by Barz et al. [1], we let $X := \{x_t\}_{t=1}^n$ be a time series for $x_t = [x_t^1, \dots, x_t^d]^\top \in \mathbb{R}^d$, and $n, d \in \mathbb{N}$. By $x^j := \{x_t^j\}_{t=1}^n$, for $j \in \mathcal{J} := \{1, \dots, d\}$, we denote each of the d variables constituting $\{x_t\}_{t=1}^n$. The set of all subintervals of the set $\{1, \dots, n\}$ is denoted by \mathcal{I} . To find the most anomalous interval $I = \{t \in \mathbb{N} \mid a \leq t < b\} = [a, b) \cap \mathbb{N} \in \mathcal{I}$, for $a, b \in \{1, \dots, n\}$, of the time series $\{x_t\}_{t=1}^n$, the MDI algorithm looks for an interval whose data distribution p_I is the most different from the distribution of the data in the rest of the time series with indices in $\Omega = \{1, \dots, n\} \setminus I$, denoted by p_Ω . This difference of the distributions p_I and p_Ω is quantified using the Kullback-Leibler (KL) divergence:

$$\mathcal{D}(p_I, p_\Omega) = \int p_I(x_t) \cdot \log \left(\frac{p_I(x_t)}{p_\Omega(x_t)} \right) dx_t. \quad (1)$$

The probability density functions $p_I = \mathcal{N}(\mu_I, S_I)$ and $p_\Omega = \mathcal{N}(\mu_\Omega, S_\Omega)$, for $\mu_I, \mu_\Omega \in \mathbb{R}^d$ and $S_I, S_\Omega \in \mathbb{R}^{d \times d}$, are approximated by a multivariate normal distribution, allowing for a closed-form solution of the KL divergence:

$$\begin{aligned} \mathcal{D}(p_I, p_\Omega) = & -\frac{1}{2} \left((\mu_\Omega - \mu_I)^\top S_\Omega^{-1} (\mu_\Omega - \mu_I) \right. \\ & \left. + \text{tr}(S_\Omega^{-1} S_I) + \log \frac{|S_\Omega|}{|S_I|} - d \right). \end{aligned} \quad (2)$$

$$\begin{aligned} \text{cov}(X') &= \mathbb{E}_t(x'_t \cdot x'_t{}^\top) = \mathbb{E}_t \left(\begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \end{bmatrix} \cdot \begin{bmatrix} x_t^\top & x_{t-1}^\top & x_{t-2}^\top \end{bmatrix} \right) = \mathbb{E}_t \left(\begin{bmatrix} x_t^1 \\ x_t^2 \\ x_{t-1}^1 \\ x_{t-1}^2 \\ x_{t-2}^1 \\ x_{t-2}^2 \end{bmatrix} \cdot \begin{bmatrix} x_t^1 & x_t^2 & x_{t-1}^1 & x_{t-1}^2 & x_{t-2}^1 & x_{t-2}^2 \end{bmatrix} \right) \\ &= \mathbb{E}_t \left(\begin{bmatrix} x_t^1 x_t^1 & x_t^1 x_t^2 & x_t^1 x_{t-1}^1 & x_t^1 x_{t-1}^2 & x_t^1 x_{t-2}^1 & x_t^1 x_{t-2}^2 \\ x_t^2 x_t^1 & x_t^2 x_t^2 & x_t^2 x_{t-1}^1 & x_t^2 x_{t-1}^2 & x_t^2 x_{t-2}^1 & x_t^2 x_{t-2}^2 \\ x_{t-1}^1 x_t^1 & x_{t-1}^1 x_t^2 & x_{t-1}^1 x_{t-1}^1 & x_{t-1}^1 x_{t-1}^2 & x_{t-1}^1 x_{t-2}^1 & x_{t-1}^1 x_{t-2}^2 \\ x_{t-1}^2 x_t^1 & x_{t-1}^2 x_t^2 & x_{t-1}^2 x_{t-1}^1 & x_{t-1}^2 x_{t-1}^2 & x_{t-1}^2 x_{t-2}^1 & x_{t-1}^2 x_{t-2}^2 \\ x_{t-2}^1 x_t^1 & x_{t-2}^1 x_t^2 & x_{t-2}^1 x_{t-1}^1 & x_{t-2}^1 x_{t-1}^2 & x_{t-2}^1 x_{t-2}^1 & x_{t-2}^1 x_{t-2}^2 \\ x_{t-2}^2 x_t^1 & x_{t-2}^2 x_t^2 & x_{t-2}^2 x_{t-1}^1 & x_{t-2}^2 x_{t-1}^2 & x_{t-2}^2 x_{t-2}^1 & x_{t-2}^2 x_{t-2}^2 \end{bmatrix} \right), \quad \text{for } t \in \{1, \dots, n\} \end{aligned}$$

Fig. 2: Covariance matrix of the three-dimensional time-delay embedding X' of a two-dimensional time series X . For simplicity, we here assume that X' has zero mean.

In order to account for the fact that the time series' samples are inter-dependent, the MDI algorithm applies time-delay embedding [8] as a pre-processing step. It creates a modified time series $X' := \{x'_t\}_{t=1+(\kappa-1)\tau}^n$, $x'_t \in \mathbb{R}^{\kappa d}$, where each sample x'_t includes attributes from previous $(\kappa - 1)\tau$ time steps as context:

$$x'_t = \begin{bmatrix} x_t^\top & x_{t-\tau}^\top & x_{t-2\tau}^\top & \dots & x_{t-(\kappa-1)\tau}^\top \end{bmatrix}^\top, \quad (3)$$

where κ denotes the number of aggregated samples called the *embedding dimension* and τ denotes the distance between two consecutive samples called the *lag*. The underlying optimization problem for finding the most anomalous interval is: $\hat{I} = \underset{I \in \mathcal{I}}{\text{argmax}} \mathcal{D}(p_I, p_{\Omega(I)})$. However, according to the user's input, the MDI algorithm can output multiple, non-intersecting, anomalous intervals with the anomaly score ranked from the highest to the lowest. Furthermore, one can configure minimum and the maximum of the intervals to consider based on data-related domain knowledge in order to improve the algorithm's performance.

IV. METHOD

Let $I = [a, b) \cap \mathbb{N}$ be an anomalous interval of the multivariate time series X , detected by the MDI algorithm. Furthermore, let x^j , $j \in \mathcal{J}$, be the variables constituting X . To attribute the anomaly that occurred in I , we employ counterfactual reasoning in the sense that we intervene on a subset of variables of X by making their values within I more like the rest of the data within $\Omega = \{1, \dots, n\} \setminus I$. We then observe if the anomaly would still have occurred had this action been taken, i.e., had a certain subset of variables been less anomalous.

More specifically, we systematically modify a subset $X_V := \{x^k \mid k \in V\}$ of variables with indices in $V \subset \mathcal{J}$, for $|V| \leq \lfloor \frac{|I|}{2} \rfloor$, by replacing their values within the anomalous interval I with a random sample from the outside distribution p_Ω . We assume that this is the nominal distribution of the data. We then re-apply the MDI scoring algorithm to the corrected interval to obtain a new anomaly score. Suppose the new score is lower than before we modified a particular subset of variables X_V . Then, if all other subsets of equally many variables yield a higher anomaly score after the same procedure, we can conclude that the variables in X_V contribute to the anomaly because of the type of the performed intervention that we call the *counterfactual replacement*.

We will now define the counterfactual replacement of a subset of variables X_V within the interval I in more detail. As mentioned previously, it answers the question of how the variables of X_V would look like inside the interval I if they would have been governed by the same distribution p_Ω as the rest of the data within Ω . Clearly, the values of the variables from X_V depend on the values of the variables from the complementary subset $X_{\bar{V}}$ in the same interval, which we do not replace. An independent sample of X_V would destroy these correlations and hence introduce another, potentially even stronger anomaly than the original one. Therefore, the replacement must preserve the correlations with the non-replaced variables. If not for the temporal dependencies between the consecutive time steps, we could simply draw $|I|$ independent samples from the conditional distribution $p_\Omega(\cdot \mid X_{\bar{V}})$ as a replacement for X_V . However, due to the time-delay embedding of the MDI algorithm, we must not only heed the inter-variable correlations but also the temporal dependencies in the data to obtain a valid counterfactual replacement. Furthermore, the

replacement should be conditional on the left and right context of the interval I , in order not to introduce an artificial anomaly due to a sudden break of temporal consistency.

In analogy to the principle of time-delay embedding, we model each time step to be replaced as an individual random variable. Following the default distribution assumption of the MDI algorithm, we then use a multivariate normal distribution $\mathcal{N}(\mu, S)$, with $\mu \in \mathbb{R}^{d\ell}$, $S \in \mathbb{R}^{d\ell \times d\ell}$, for estimating the joint distribution of all subsets of the time series of length $\ell = b - a + 2(\kappa - 1)$. This corresponds to the length of the anomalous interval plus the left and right context window, whose length depends on the time-delay embedding dimension κ used by MDI for anomaly detection. Obtaining the mean μ of this distribution is straightforward: Since it does not depend on the position of the samples in the time series, we can simply repeat the mean μ_Ω of the nominal distribution ℓ times.

The covariance matrix S , in contrast, would be quite large and difficult to estimate robustly from limited data. We observe, however, that in the limit $n \rightarrow \infty$ for infinitely long time series, S converges to a symmetric block Toeplitz matrix, i.e., the blocks on each diagonal are constant. The small illustrative example in Fig. 2 exhibits this property for a two-dimensional time series and $\ell = 3$. The blocks having the same color code in Fig. 2 have the same expectation value in the limit. Therefore, we only need to estimate the first row of blocks from the available data, which fully defines the entire covariance matrix. This reduces the computational complexity drastically from $\mathcal{O}(\ell^2 d^2)$ to $\mathcal{O}(\ell d^2)$. To prevent the anomaly itself from influencing the estimate of the nominal distribution, we replace the interval I with missing values before computing μ and S .

Finally, we draw a replacement for the selected variables X_V within the anomalous interval I from the conditional distribution $\mathcal{N}(\mu, S \mid X_{\bar{V}}, X_{\text{left}}, X_{\text{right}})$. We condition on the non-replaced variables $X_{\bar{V}}$ with $\bar{V} = \mathcal{J} \setminus V$ to maintain the inter-variable dependencies and on the left and right context $X_{\text{left}} = [x_{a-(\kappa-1)}, \dots, x_{a-1}]^\top$, $X_{\text{right}} = [x_b, \dots, x_{b+(\kappa-2)}]^\top$ to make the replacement connect smoothly to the surroundings of the anomalous interval. An example of such a counterfactual replacement is shown in Fig. 1.

V. EXPERIMENTS

We applied our method to several well-understood extreme weather events and compared the attributed variables as the means of evaluation. In all the experiments, we used the time-delay embedding of dimension $\kappa = 3$, the lag $\tau = 1$ and the unbiased KL divergence $\text{U-KL}(p_I, p_\Omega) := 2 \cdot |I| \cdot \text{KL}(p_I, p_\Omega)$ as rationalized by Barz et al. [1]. Furthermore, we specified the size of the desired anomalous interval for each dataset individually. We applied our method to three different datasets of multivariate time series and one multivariate spatio-temporal dataset. The results in Tables I-IV are obtained after averaging 10 realizations of the in-distribution replacements of our attribution method scored by the KL divergence. In all our experiments, we discuss the attribution during and before the detected anomalous intervals to account for the lagged-effect of certain variables, i.e., whether the anomalous interval is an

		10 days before	Detection
Anomaly score		104.69	371.44
SLP		93.91	332.97
	W	79.09	350.25
	Hs	105.58	248.34
SLP	W	57.11	333.4
SLP	Hs	107.32	129.7
	W Hs	61.81	182.58

TABLE I: Anomaly scores after counterfactual replacements on the Hurricanes dataset. The lowest scores for variable subsets of different cardinality are highlighted and shown in bold font.

effect of a cause that started earlier in time than the detected interval.

Our source code is available at <https://github.com/cvjena/mdi-attribution>.

A. Hurricanes

We first apply our attribution approach to meteocean data obtained near the Bahamas in the Atlantic Sea from the National Data Buoy Center from the National Oceanic and Atmospheric Administration (NOAA) [9]. It encompasses six months of hourly measurements of significant wave height (Hs), wind speed (W), and sea level pressure (SLP), from June 2012 until November 2012. During this time, the Atlantic hurricane season was especially active with the occurrence of 19 tropical cyclones ($W > 52$ km/h), 10 of which became hurricanes ($W > 64$ km/h).

We set the size of the intervals to be searched for by the MDI algorithm between 24 and 120 hours. In Fig. 1, we show the interval replacement of the Hs variable within the detected anomalous interval. Table I shows our attribution results during the detected hurricane, as well as ten days before the event. The variable names indicate which subset of variables was counterfactually replaced within the anomalous interval I . When only one variable at a time is replaced within I , we observe that the anomaly score improves the most in the case of the wave height. This means that the MDI algorithm mostly detects this interval as anomalous due to the anomaly in Hs. However, when inspecting the interval ten days prior to the occurrence of the detected hurricane, we notice that modifying the wind speed during that period lowers the anomaly score. When two-variable subsets are counterfactually replaced within I , we note that W helps lower the anomaly score even further during the event and can therefore be attributed to the MDI's detection together with Hs. However, ten days prior to the event, SLP along with W reduces the anomaly score the most when their values within I are replaced by the in-distribution samples. This could imply that SLP is the cause of the hurricane itself. As expected, when we repair all three variables, the anomaly score drastically decreases. For this reason, in all the experiments we assume that at most $\frac{d}{2}$ variables are responsible for the event. Otherwise, by changing more variables in the interval I , we drastically change the settings of the event which could result in an inaccurate attribution.

		DE-Hai	One month before	Detection
		Anomaly score	21.63	142.57
T_{air}	NEE		20.54	140.99
		PPT	19.7	133.89
		VPD	21.16	51.17
		LE	20.72	113.45
		H	20.52	141.03
T_{air}	NEE		19.1	135.04
T_{air}		PPT	19.64	115.53
T_{air}		VPD	18.8	37.39
T_{air}		LE	19.59	99.51
T_{air}		H	19.63	136.95
T_{air}	NEE	PPT	19.4	112.61
	NEE	VPD	20.34	50.27
	NEE	LE	18.08	109.38
	NEE	H	20.75	134.87
	PPT	VPD	20.26	30.13
	PPT	LE	18.98	90.85
	PPT	H	18.95	118.53
	VPD	LE	20.34	46.29
	VPD	H	20.01	44.63
	LE	H	20.38	109.95
T_{air}	NEE	PPT	18.01	111.09
T_{air}	NEE	VPD	18.37	37.12
T_{air}	NEE	LE	16.07	87.58
T_{air}	NEE	H	20.39	135.07
T_{air}	PPT	VPD	18.35	18.6
T_{air}	PPT	LE	17.06	79.11
T_{air}	PPT	H	17.58	113.83
T_{air}	VPD	LE	16.74	36.36
T_{air}	VPD	H	19.07	34.45
T_{air}	LE	H	18.07	99.87
T_{air}	NEE	PPT	18.37	31.24
T_{air}	NEE	VPD	16.98	85.75
T_{air}	NEE	LE	19.25	114.6
T_{air}	NEE	H	18.39	44.17
T_{air}	VPD	LE	19.26	43.34
T_{air}	VPD	H	17.74	102.92
T_{air}	PPT	VPD	20.35	26.2
T_{air}	PPT	LE	18.1	24.75
T_{air}	PPT	H	18.4	85.02
T_{air}	VPD	LE	19.32	42.09

TABLE II: Anomaly scores after the counterfactual replacements one month before and during the detected anomaly from July and August 2003 on the German site DE-Hai. The lowest scores for variable subsets of different cardinality are highlighted and shown in bold font.

B. Multivariate Ecological Time Series

We perform anomaly attribution in multivariate time series of the certain ecological sites of the FLUXNET [10] dataset. Particularly, we analyze the causes of the European heatwave in the Summer of 2003 in Germany and France, by using the data from the site DE-Hai and FR-Pue, respectively.

1) *German Heatwave Attribution:* To investigate the extreme precipitation event, namely the lack of rain that occurred in Germany in 2003 during the European heatwave, we take a closer look at the FLUXNET site DE-Hai, located in the Hainich National Park. We use the daily measurements of the air temperature (T_{air}), net ecosystem exchange (NEE), precipitation (PPT), vapor pressure deficit (VPD), latent heat flux (LE), and sensible heat (H), as suggested by Krich et al. [11]. The data span from the year 2000 until 2009. We first apply the MDI method to detect the two-month-long anomalous interval, which coincidentally happens to occur exactly from July until the end of August 2003, when the European heatwave had occurred as well. Then we perform systematic in-distribution replacements within the detected interval for subsets of variables with at most three variables and re-score it. The results thereof are shown in Table II. When considering singleton variable subsets, we note that the most responsible variable for the MDI’s detection of this interval is VPD. When replacing an anomalous segment of

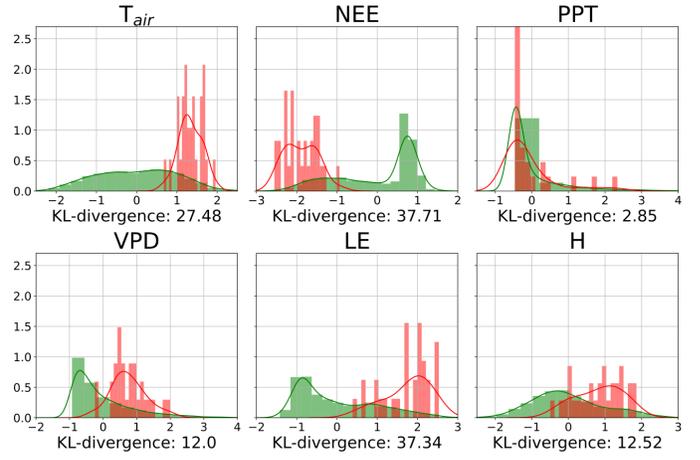


Fig. 3: Attribution scheme of the German heatwave. The histogram of the entire variable is shown in green, whereas the histogram of the variable within the anomalous interval is shown in red. The KL-divergence results of the univariate comparison of the red and green histograms are shown individually for each variable under the x -axis.

two variables at a time, we find that this extreme event can be attributed to PPT, along with VPD, which is in line with Krich et al. [11]. Moreover, when we replace three variables at a time within the detected interval, we observe that T_{air} in addition to PPT and VPD reduces the anomaly score the most. This is also a well-supported attribution when studying an extreme event such as a heatwave. In addition, we inspected the one-month-long interval before the detected event and observed that the most anomalous singleton consists of PPT, the most anomalous two-variable subset consists of NEE and LE and adding T_{air} to the latter two variables seemed to have been the most dissimilar to the rest of the data one month prior to the event. This may have contributed to the occurrence of the said heatwave.

To emphasize the importance of considering inter-variable correlations in multivariate time series during attribution, we furthermore conduct a simple baseline attribution experiment: Each variable is considered in isolation and the anomaly scores during the time frame in question are compared. Fig. 3 shows that such a purely univariate approach would arrive at the false conclusion of NEE and LE being the most anomalous variables, whereas we know from our multivariate attribution scheme that VPD and PPT actually account for the anomaly.

2) *French Heatwave Attribution:* The FLUXNET site FR-Pue which we used to inspect the European heatwave in France, consists of the 14 years of daily measurements of the ecological variables T_{air} , NEE, PPT, VPD, LE, and H, in the period from 2001 until 2014. In Table III, we see the results of counterfactually replacing different subsets of variables within the anomalous interval, while carefully minding the inter-temporal relations. We note that the anomaly was detected from mid-July until mid-August of 2003, whereas the heatwave was documented through the entire July and August of the

FR-Pue				One month before	Detection
Anomaly score				20.93	123.38
T_{air}	NEE	PPT	VPD	18.9	113.78
				17.61	116.88
	LE	H	20.56	83.61	
			18.77	83.54	
			17.9	120.91	
20.88	121.26				
T_{air}	NEE	PPT	VPD	16.37	110.21
				17.83	75.35
	LE	H	17.79	74.02	
			14.5	109.76	
			18.61	110.51	
	NEE	PPT	VPD	18.66	77.07
				17.19	73.74
	NEE	VPD	LE	14.99	97.43
				18.44	115.41
	NEE	PPT	VPD	17.56	42.88
15.99				81.86	
PPT	VPD	LE	20.69	82.5	
			16.88	86.95	
VPD	LE	H	19.66	79.58	
			19.05	124.36	
T_{air}	NEE	PPT	VPD	18.24	69.55
				17.35	67.88
	NEE	VPD	LE	14.43	88.04
				18.21	107.98
	PPT	VPD	LE	14.73	35.28
				13.94	69.98
	PPT	VPD	LE	18.27	70.36
				13.84	82.69
	VPD	LE	H	18.17	61.46
				14.7	106.02
NEE	PPT	VPD	16.51	35.58	
			14.46	55.12	
NEE	PPT	VPD	18.71	75.73	
			14.8	65.91	
NEE	VPD	LE	16.91	71.76	
			16.55	96.49	
NEE	VPD	LE	15.5	47.42	
			18.66	39.89	
PPT	VPD	LE	17.88	78.73	
			17.52	82.83	

TABLE III: Anomaly scores after counterfactual replacements one month before and during the detected anomaly from mid-July to mid-August 2003 on the French site FR-Pue. The lowest scores for variable subsets of different cardinality are highlighted and shown in bold font.

same year. When counterfactually modifying only one variable at the time, we see in Table III that VPD lowers the anomaly score the most, meaning that it also contributes to the anomaly the most. In the case of substituting the anomalous segments of two variables at the time, PPT and VPD yield the lowest anomaly score, and we attribute this heatwave to them. This is similar to the attribution of the same event on the German site DE-Hai. When replacing three variables' values with the in-distribution ones within the detected interval, we note that T_{air} contributes to lowering the anomaly score in addition to PPT and VPD. This is in line with Shadaydeh et al. [6] who also attributed this event to T_{air} and VPD. Furthermore, as for the German site, we analyzed the one-month-long interval prior to the detected anomaly and noted that in this case NEE was the most anomalous, potentially contributing to the occurrence of the heatwave in question. Two-variable replacements a month before the detection indicated that the event could have been caused by abnormalities of T_{air} and LE. Moreover, when simultaneously replacing three variables with the in-distribution values one month prior to the detected anomaly, we see that T_{air} , VPD and LE yield the lowest anomaly score.

C. Spatio-Temporal Data

For the experiments on spatio-temporal data, we use the data provided by Racah et al. [12]. It consists of spatio-temporal

Spatio-temporal data		10 days before	Detection
Anomaly score		4.67	6.11
SLP		200	340
TS		60	2.69
	W	90	380
SLP	TS	2.6	3.1
SLP	W	10	50
	TS	6.71	2.95

TABLE IV: Anomaly scores after counterfactual replacements 10 days before and during the 12-day-long detected extreme weather event in spatio-temporal data. The lowest scores lower than the anomaly score, for variable subsets of different cardinality are highlighted and shown in bold font. All values are given in units of 10^5 .

variables of the weather conditions on Earth and boxes outlying the hurricanes along with a class label. Four samples per day are available for a period of 365 days. For our anomaly attribution analysis, we choose the variables of sea level pressure (SLP), surface temperature (TS), and wind speed (W), similarly to the variables of the purely temporal Hurricanes dataset.

We apply our method pixel-wise to a region of 220×220 pixels, where we know a hurricane has occurred, to achieve computational tractability. We preset the spatial size of the detection for the spatio-temporal MDI algorithm to regions of 50×50 pixels, lasting 12 days. In Table IV, we show our method's results ten days before and during the detected extreme event. The detected event can be attributed to TS as counterfactually replacing that variable within the detection interval reduces the anomaly score the most. Interestingly, when we counterfactually replace two variables at a time within the corresponding time frames, in addition to TS, ten days before the anomaly, SLP influences the anomaly score the strongest. In contrast, during the detection itself, the anomaly score is influenced the most when replacing W and TS.

VI. CONCLUSION

We proposed a novel anomaly attribution approach for multivariate temporal and spatio-temporal data based on the MDI method for anomaly detection and the counterfactual replacements of variables within the anomalous intervals. This counterfactual replacement answers the question of would the anomaly still have occurred, had a subset of variables been more similar to the data outside of that interval. The main benefit of our attribution method is that it can be applied to any multivariate time series data, regardless of potential outliers or missing values, since it replaces the entire anomalous interval with in-distribution samples and re-scores it. It also considers the correlations both among variables and across time. It is thus also suitable for the task of data repairing. Furthermore, it is one of the few methods for multivariate anomaly attribution in time series, and provides interpretation of the MDI method's anomaly detections. We demonstrated the use of our method on multiple multivariate time series datasets and one multivariate

spatio-temporal dataset. By obtaining attributions for well-documented extreme events whose causes are well-understood, we confirmed the correctness of our method. Therefore, we conclude that our approach is a fast and accurate tool to help domain experts in understanding the possible causes of anomalous events in multivariate time series. Moreover, the current work can be further exploited towards the classification of different anomalous events based on the presented attribution.

ACKNOWLEDGMENT

This work is funded by the European Research Council (ERC) Synergy Grant 2019: Understanding and Modelling of the Earth System with Machine Learning (USMILE) and by the Carl Zeiss Foundation within the scope of the program line "Breakthroughs: Exploring Intelligent Systems" for "Digitization — explore the basics, use applications".

REFERENCES

- [1] B. Barz, E. Rodner, Y. Guancho Garcia, and J. Denzler, "Detecting regions of maximal divergence for spatio-temporal anomaly detection," vol. 41, pp. 1088–1101, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [2] A. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, "A review on outlier/anomaly detection in time series data," *ACM Comput. Surv.*, vol. 54, no. 3, pp. 1–33, 2021.
- [3] M. S. A. Siddiqui, D. Mercier, A. Dengel, and S. Ahmed, "Tsinsight: A local-global attribution framework for interpretability in time-series data," *ArXiv e-prints*, vol. abs/2004.02958, pp. 1–16, 4 2020.
- [4] T. Idé, A. Dhurandhar, J. Navrátil, M. Singh, and N. Abe, "Anomaly attribution with likelihood compensation," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 5, pp. 4131–4138, 2021.
- [5] A. Zhang, S. Song, J. Wang, and P. S. Yu, "Time series data cleaning: From anomaly detection to anomaly repairing," vol. 10, *Proceedings of the VLDB Endowment*, 2017.
- [6] M. Shadaydeh, J. Denzler, Y. Guancho Garcia, and M. Mahecha, "Time-frequency causal inference uncovers anomalous events in environmental systems," in *German Conference on Pattern Recognition (GCPR)*, pp. 499–512, 2019.
- [7] Y. Guancho Garcia, M. Shadaydeh, M. Mahecha, and J. Denzler, "Attribution of multivariate extreme events," in *International Workshop on Climate Informatics (CI)*, 2019.
- [8] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw, "Geometry from a time series," in *Physical review letters*, vol. 45, pp. 712–715, 1980.
- [9] B. Barz, Y. Guancho, E. Rodner, and J. Denzler, "Maximally divergent intervals for extreme weather event detection," in *MTS/IEEE OCEANS Conference Aberdeen*, pp. 1–9, 2017.
- [10] D. Baldocchi, E. Falge, L. H. Gu, *et al.*, "Fluxnet: A new tool to study the temporal and spatial variability of ecosystem-scale carbon dioxide, water vapor, and energy flux densities," *Bulletin of the American Meteorological Society*, vol. 82, no. 11, pp. 2415–2434, 2001.
- [11] C. Krich, M. Migliavacca, D. G. Miralles, G. Kraemer, T. S. El-Madany, M. Reichstein, J. Runge, and M. D. Mahecha, "Functional convergence of biosphere-atmosphere interactions in response to meteorological conditions," *Biogeosciences*, vol. 18, no. 7, pp. 2379–2404, 2021.
- [12] E. Racah, C. Beckham, T. Maharaj, S. Kahou, M. Prabhat, and C. Pal, "Extremeweather: A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events," in *Advances in Neural Information Processing Systems 30* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), pp. 3405–3416, Curran Associates, Inc., 2017.