Oxford Open Climate Change, 2024, **4(1)**, kgae009
https://doi.org/10.1093/oxfclm/kgae009
Advance Access Publication Date: 28 June 2024
Research Article

Identifying when thresholds from the Paris Agreement are breached: the minmax average, a novel smoothing approach

Mathieu Van Vyve [1,2,*]

[1]Center for Operations Research and Econometrics, LIDAM, UCLouvain, voie du Roman Pays 34, Louvain-la-Neuve, Wallonia, 1348, Belgium
[2]Sauder School of Business, University of British Columbia, University of British Columbia, Vancouver, British Columbia, V6T 1Z2, Canada

*Corresponding author. Center for Operations Research and Econometrics, LIDAM, UCLouvain, Voie du Roman Pays 34, 1348, Belgium. E-mail: mathieu.vanvyve@uclouvain.be

Abstract

Identifying when a given threshold has been breached in the global temperature record has become of crucial importance since the Paris Agreement. However there is no formally agreed methodology for this. In this work we show why local smoothing methodologies like the moving average and other climate modeling based approaches are fundamentally ill-suited for this specific purpose, and propose a better one, that we call the minmax average. It has strong links with the isotonic regression, is conceptually simple and is arguably closer to the intuitive meaning of "breaching the threshold" in the climate discourse, all favorable features for acceptability. When applied to the global mean surface temperature anomaly (GMSTA) record from Berkeley Earth, we obtain the following conclusions. First, the rate of increase has been ∼+0.25°C per decade since 1995. Second, based on this new estimate alone, we should plausibly expect the GMSTA to reach 1.49°C in 2023 and not go below that on average in the medium-term future. When taking into account the record temperatures of the second half of 2023, not having breached the 1.5°C threshold already in July 2023 is only possible with record long and/or deep La Niña in the following years.

Keywords: Paris Agreement; threshold Identification; smoothing; isotonic regression

Identifying if and when the thresholds of 1.5°C or 2°C are breached is key for policy

The temperature on Earth is very variable in space and time, but it is now widely recognized that it is gradually and on average increasing and that this constitutes a major challenge to the human civilization. Compared to the pre-industrial period, the global mean surface temperature (GMST) is now well above 1°C higher, and this already has major negative impacts on human societies through more frequent and severe extreme weather events, glacier loss and stress for many plants, among others. Substantially even higher temperatures could have catastrophic long-term consequences for billion of humans and other species, through generalized crop failures, multiple-meters sea level rise, lack of access to fresh water, and large-scale migration [1–4].

Fortunately this danger has been recognized by the leaders of our societies, and it is now part of an international treaty, the 2015 Paris Agreement, that humanity will take necessary action to keep GSMT increase well below 2°C compared to pre-industrial levels, with the aim of keeping it below 1.5°C. Because the cause of the warming of the climate is mostly man-made, human societies are usually evolving slowly, and the rate of increase during the last decades has been at least of 0.5°C per 30 years, this is a formidable task [5]. Nonetheless, agreeing on these objectives was a huge step for policy reasons: approaching these limits, in particular when faster than anticipated, or altogether breaching them, is a clear signal that more needs to be done.

Maybe surprisingly given the stakes, there is no formally agreed methodology to determine when a certain threshold is reached. There are at least three difficulties of very different nature here. The first one is to determine the baseline with which to compare the current GMST. It must be sufficiently early to pre-date the large-scale usage of fossil fuels, while also guaranteeing the good quality of the available temperature measurements. The IPCC has resolved this tradeoff by using as the baseline the average temperature in the period 1850–1900 [6]. On the one hand, some authors have argued the baseline should be earlier, because there has been a non-negligible temperature increase between 1750 and 1850 [7]. On the other hand, there is still an active debate about the actual temperature in 1850–1900 as the most widely used datasets differ by up to 0.2°C for that period [8]. These questions are not the main topic of this work, and we have chosen to use the Berkeley Earth Data series with the IPCC baseline of 1850–1900. Other choices would yield similar conclusions, simply offset by a constant value (e.g. the so-called structural uncertainty), because there is good agreement about the global temperature in the modern era [8].

The two other difficulties are both related to the fact that the global temperature of Earth as a function of time is not monotone: it is going up and down. For example, it is widely recognized that the global temperature anomaly went above 1°C for several

Received: March 07, 2024. **Revised:** June 03, 2024. **Accepted:** June 05, 2024

months in the period 1997–1999, but this was followed by consistently lower temperature in 2000 and 2001, and that is why it is widely accepted that the 1°C threshold was actually reached many years later. Conversely, it does not make sense to wait for the temperature anomaly to *never* be under 1.5°C for 20 years to consider the threshold breached. Once consistently above, a temporary dip below 1.5°C would not be a sign that humanity is doing well. Clearly, some averaging or smoothing over time is necessary. But how exactly? Answering this question is the main topic of the present work.

A third difficulty is that whether a given threshold is attained today clearly depends on what will happen in the next, future, years. For example, in its latest report AR6, the IPCC made a specific choice: to use a 20-years moving average centered around the year of interest [6]. This implies that, in principle, one can only know the quantity of interest with a 10-years delay, and that is catastrophically late for policy, which is precisely the intended use of this indicator. Clearly, a similar difficulty will exist for any sensible averaging method. Betts *et al*. recently discussed this issue and meaningfully propose to use simulated data for future years, in order to be able to estimate whether a given threshold is reached without delay [9]. Using simulated data will introduce additional uncertainty in the estimate, but this is much better than nothing, and that uncertainty itself can usually be estimated to inform policy.

Therefore, we will in general assume here that a sufficiently long time series T_i is available, giving the temperature (or temperature anomaly) for each time period i in an interval of integers ranging from M to N. Such time periods might be days, months or years, and could span the past and/or the future. This data can originate from measurement or simulation. Our goal in this note is to develop a sound methodology to determine in which period i in the series a given threshold L is reached. Only when applying our methodology to very recent climate data will we discuss how to deal, if possible, with the uncertainty regarding the future.

The rest of this note is organized as follows. In the next (second) section, we discuss why current proposals are not suitable for determining when a threshold is attained in a time series. In the third section we discuss the advantages and drawbacks of using the isotonic regression instead. The fourth section describes our new minmax average approach, closely linked to the isotonic regression. In the fifth section we apply our methodology to synthetic data, the temperature anomaly past record, and the determination of the minmax average in 2023. We finally discuss potential issues with our proposed methodology in the sixth section before concluding. For notational convenience, in the rest of this paper, we will denote the average temperature within an interval of time periods $[i,j]$ by $T_{i,j} = \frac{\sum_{k=i}^{j} T_k}{j-i+1}$. Also, we are mostly interested in applying our methodology to the GMSTA record, but it could in principle be applied to any time series.

Issues with currently proposed methods

Local smoothing

As already mentioned, there is a default methodology that has been used by the IPCC in AR6: a 20-years moving average centered at the period of interest (note that a longer, 30-years averaging period, is used by the WMO and IPCC in Assessment Reports up to AR5, but that has little impact for the discussion here). As clearly illustrated in Fig. 1, this satisfies the goal of averaging ups and downs with short cycles, and the resulting smoother curve is indeed making the trend of the time series clearly visible to the untrained eye. The answer to what we will call here the *threshold question* "When is a certain threshold L reached in the time series T?", one simply then needs to identify the first period where this value L is attained in the 20-years moving average.

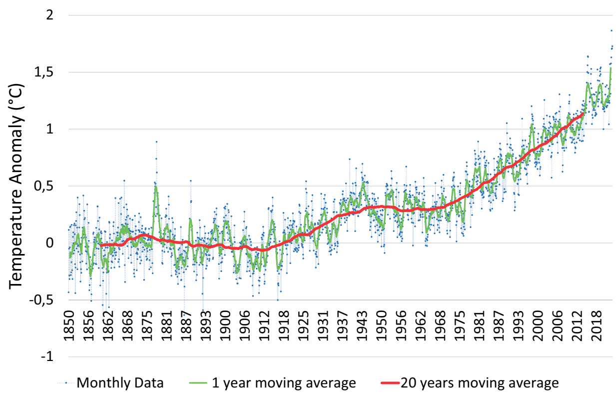

Figure 1. Temperature anomaly from period 1850–1900, 1- and 20-years centered moving averages, Berkeley Earth Data [8].

Betts *et al*. give another argument in favour of this simple method [9]. As they put it: "Any definition must be consistent with how 1.5°C is already defined by the IPCC; that is, using 20-year averages attached to a midpoint." Although we agree that ensuring methodological stability over time is important for credibility, this cannot be an absolute scientific principle. If we can show that this method suffers from basic and obvious flaws, this would also be bad for credibility. The better option might very well then be to change the methodology to maximize credibility, all things considered. At a minimum, having multiple methodologies at hand with identified weaknesses is more helpful than a single flawed one.

And indeed, there exist specific cases where this centered moving-average approach will give results that do not make sense. The simplest such case is when the temperature time series follows a step increase from one level to another higher one, as illustrated in Fig. 2, example I. Note first that such a one-off increase is not a purely hypothetical scenario. For example the very rapid desulphurization of marine fuels is expected to qualitatively have such an impact [10]. Using the moving average approach, one would conclude that the 1°C threshold is reached in 2019. But this is clearly wrong: it suffices to read on the original, non-smoothed, temperature time series that it reaches 1°C in 2009, 10 years before. Conversely, using the moving-average approach, one would conclude that the 0.6°C threshold is attained in 2002. But again, this statement does not make sense, since this temperature level is never reached before 2009. Also it is quite clear that more sophisticated but similar local smoothing methodologies, like LOESS, will suffer from the same basic problem [11].

In general, the 20-years moving-average approach will only leave linear time series unchanged, and all other type of functions/series will be modified. But this discussion shows that no smoothing at all should be applied as long as the time series is monotonously increasing. Or equivalently, a reasonable averaging methodology, when applied to any monotone series/function, should naturally (i.e. without treating it as a special case) leave it unchanged. Only ups *and* downs should be smoothed out.

Model-based approaches

A strategy that has been advocated is to leverage detection and attribution studies to remove from the temperature signal some

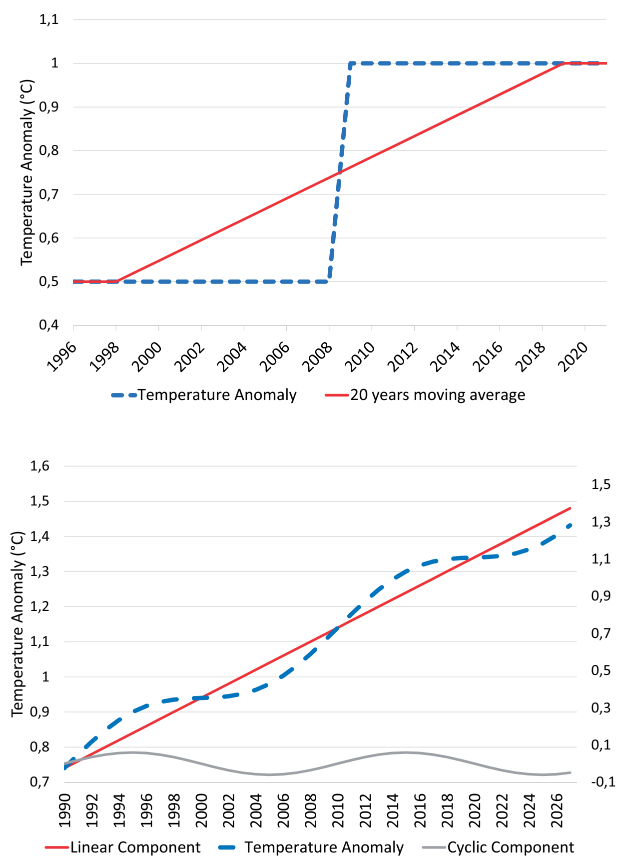

Figure 2. Two stylized examples of monotone temperature profiles. Example I is a one-off step increase and the 20-years moving average. Example II is the sum of two components: a linear one and a cyclic one.

natural, high-frequency cyclic or noisy components to obtain a human-induced temperature anomaly time series, which yields an index that is much smoother and monotone, as the real-time Global Warming Index [12–16]. This could be used as a smoothing strategy to answer the threshold question. Diffenbaugh *et al.* also use climate models combined with machine learning to remove cyclic components from the temperature record [17].

These approaches are insightful, but suffer from similar problems as the moving average, as illustrated by the fictive situation depicted in Fig. 2, Example II, where the global temperature is the result of the addition of two components: the first one gradually and monotonously increasing, and the second one cyclic with mean zero. In that case these approaches would filter out the second component and extract the first one. Based on this, we would for example conclude that the 1°C mark has been reached in 2003 in the example. But that would, with reason, lack credibility, as it is clear for everyone that never before 2006 was this temperature observed.

An implicit assumption in the reasoning above is that it is the actual temperature anomaly that is the object of the Paris Agreement, and not a (model-based) index stripped from its natural, cyclic or transient components. This interpretation is supported by the text itself, Article 2.1.a: "Holding the increase in the global average temperature to well below 2°C above pre-industrial levels [...]". In addition, the ultimate aim of the agreement is to preserve the well-being of human societies, insofar as it is influenced by global temperature. The causes of the evolution of the temperature are irrelevant in that matter. If the origins of GHG (or more generally, of global temperature change) were substantially natural, the unmitigated impacts for humanity would be the same, only the possible actions and responses would be (vastly) different. Identifying the causes of climate change is certainly key to simulate and forecast the future, and to analyze, study and propose what actions could or should be taken, but it is much harder to see why that is relevant when trying specifically to answer the threshold question. Anyway, it is quite clear that a conceptually sound methodology that does not implicitly or explicitly attribute global warming to specific causes, and sticks as close as possible to the raw data, would be useful as a complement to attribution methodologies.

A first attempt at a better proposal: the monotone or isotonic regression

The discussions from the previous section naturally lead to propose the usage of a classical statistical approach known as monotone or isotonic regression. This method has characteristics of both parametric and non-parametric approaches. In the spirit of parametric methods, it asks to find the least-square approximation to the data points, but the structure imposed is non-parametric. More precisely, the isotonic regression is by definition the least-square approximation that is monotonously increasing (or decreasing, but that choice is made ex ante) [18–20]. In our setting, this is the optimal solution to the following optimization problem in a vector of variables x defined over the same set as our data T:

$$\min_{x} \quad \sum_{i=M}^{N} (x_i - T_i)^2 \qquad (1)$$

$$\text{s.t.} \quad x_i \geq x_{i-1} \qquad \forall i > M \qquad (2)$$

This convex optimization problem can be solved very efficiently, in $\mathcal{O}(n)$ computing time [21]. This approach is used nowadays in various application fields, for example in pharmacology, biology or machine learning, when the relationship is known ex ante to be increasing for structural reasons, but could potentially be substantially non-linear or non-convex [22–26]. Our context is different, because there is no direct causality between time and temperature anomaly from which we could infer monotonicity. Rather, as we already discussed above, the threshold question might not have a clear answer for some threshold L exactly because and when the data series is not monotone. In that case, it seems reasonable to use the isotonic regression instead, i.e. to compute the closest function or time series for which the question has a clear and natural answer.

As illustrated in Fig. 3 where we apply the isotonic regression to the GMSTA record, it is well known that the optimal solution x^* to problem (1)–(2) will always implicitly decompose the time interval $[M,N]$ into subintervals $\cup_{j\in[1,J]}[m_j,n_j]$ satisfying the following properties [21]:

a) Interval decomposition: $m_1 = M$, $NJ = N$ and $m_j = n_{j-1}+1$ for $j = 2,\ldots,J$.
b) The optimal solution is constant within an interval: $x^*_i = x^*_{i+1}$, except if $i = nj$ for some j.
c) The values across intervals are strictly increasing: $x^*_{m_j} < x^*_{m_j+1}$ for $j = 1,\ldots,J-1$
d) Within an interval j, the optimal value x^* is the mean of the data T in that interval: $x_{m_j} = T_{m_j,n_j}$.
e) Within an interval, the data T is decreasing, in the following sense: for any j, p such that $p \in [m_j, n_j - 1]$, $T_{m_j,p} \geq T_{p+1,n_j}$.
f) For any j,

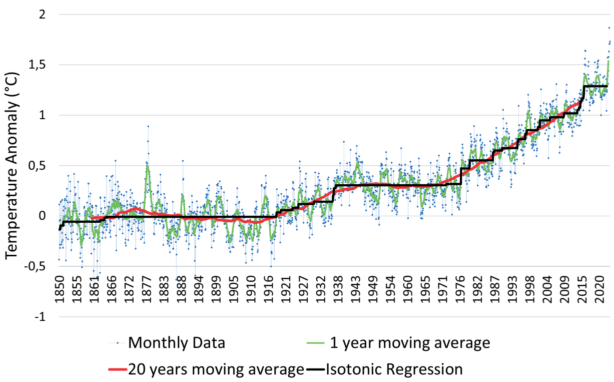

Figure 3. Temperature anomaly from period 1850–1900, 1- and 20-years moving averages and the isotonic regression, Berkeley Earth Data [8].

$$T_{m_j,n_j} = \min_{p:p\geq m_j} T_{m_j,p}. \tag{3}$$

An intuitive interpretation of (e) is that within each interval, the temperature is globally decreasing, but since we constrain the approximation x to be non-decreasing, the best possible is to keep the approximation x constant.

Let us now compare the isotonic regression and the 20-years moving average. One clear advantage of the former is that, on the condition that the moving average is monotonously increasing, as in the period 1972–2023, the isotonic regression is closer to the data, or equivalently, the least-square error is smaller. This is by construction, since both approximations are monotone and thus feasible solutions to problem (1)–(2), and the isotonic regression is the optimal one. That implies that the isotonic regression is actually tracking the data more closely, it reflects the original data better. In particular, this means that if there is a jump or a rapid increase in the isotonic regression, this corresponds to a rapid increase in the data as well. This is for example obvious in 1977, where the 1-year averages were nearly never reaching 0.45 before (and were much lower on average), while this is exactly the opposite after 1977. There is a sudden increase in the raw data around 1977, and that is not reflected in the 20-years moving average. Conversely, if the data is increasing fairly gradually, the isotonic regression would do so as well, as is the case in the period 1918–1937 for example.

Another nice consequence of being able to track rapid increases is that the isotonic regression does not "anticipate" the data. For example, the 20 years moving average markedly increases in the period 1967–1977, but this has much less to do with the temperature anomaly in that period, and more with the rapid increase post 1977.

However, the isotonic regression has the drawback of not being able to track both increases and decreases in the data, even if these are obvious 30 + years subtrends, as in 1873–1911. This is by construction, as the method requires to choose a priori whether the regression will decrease or increase.

This reasoning could suggest several adaptations to the isotonic regression approach to be able to overcome this drawback. For example, one could define a two-stage approach where first the time horizon $[M, N]$ would be partitioned into time intervals where the data is alternatively increasing or decreasing (for example by looking at the 20-years moving average) and then in a second step a corresponding optimization problem similar to (1)–(2) would be solved. Another option would be to construct a one-stage (combinatorial) optimization problem where the solution is allowed to switch from an increasing regime to a decreasing one, but only after a time interval with constant value of at least K time periods (with say $K = 20$ years), or only a certain number of times.

However, we believe these approaches all suffer in turn from other issues. And more importantly, there exists an alternative approach that is both simpler, more direct and easier to understand in the context of the threshold question that motivates this work.

Minmax average: necessary and sufficient conditions instead of a regression

To motivate the alternative proposal below, it is interesting to take a step back and discuss the intuitive, colloquial, English meaning of "breaching the threshold L at period i" in the climate context. Clearly, a necessary condition must be that $T_i \geq L$ (the temperature is higher than L in i), otherwise it would make more sense to say that the threshold was reached before or after i, but not precisely in period i. Also, this is not sufficient. In particular, if $T_i \geq L$ but $T_{i+1} << L$ so that the average temperature in the time interval $[i, i+1]$ is below L, we understand this as just a temporary overshoot in i that does not count. The low temperature $i+1$ more than compensates the high one in i. Therefore an additional necessary condition for breaching the threshold L in i must be that $T_{i,i+1} \geq L$. Now the same exact argument will apply also to the intervals $[i, i+2], [i, i+3], \ldots$ up to some specified minimum period of time K that we deem sufficient to truly conclude that the threshold has been attained, so up to $[i, i+K-1]$.

From this discussion we can conclude that, what we precisely mean by "breaching the threshold of L in period i" (for a meaningful duration of K periods), is that the following condition holds:

$$T_{i,p} \geq L \qquad \forall p \in [i, i+K-1] \tag{4}$$

or equivalently that

$$\underline{T}i^K \geq L, \text{ where } \underline{T}i^K \triangleq \min_{p\in[i,i+K-1]}. \tag{5}$$

This translates mathematically the fact that the temperature is higher than L in i, but also that it will never meaningfully (on average) go back again below L from period i to any horizon of K periods or less (i.e. in the short or medium term). Here, $\underline{T}i^K$ is the minimum temperature that will be observed on average from period i to any future period, within a horizon of K periods.

Note that these conditions are both necessary and sufficient. If any of those K conditions is not satisfied, then a good argument can be made that the threshold has not been reached. Conversely, if all of them are satisfied for sufficiently high K, it is clear that the threshold L has been reached in period i. Now, if Conditions (4) or (5) are necessary and sufficient, it logically follows that the proposed concept is the only acceptable one.

Moreover, and somewhat crucially, when $K = \infty$, Equation (5) used to define $\underline{T}i^\infty$ is exactly identical to Equation (3) satisfied by the isotonic regression. This implies that, for answering the threshold question, using the conditions (4) above with $K = \infty$ or the isotonic regression will yield the same answer and are equivalent. This shows that Conditions (4) or (5) for finite K are similar, but weaker, than checking if the isotonic regression is higher than L, because we do not average until the end of time: we only take averages up to K periods, which has more sense in the context of climate science. However, apart from that, the two approaches are structurally very similar.

Another consequence is that using Conditions (4) or (5) is equivalent to using the isotonic regression when the latter does not have constant intervals longer than K periods, or equivalently using conditions (e), as long as the temperature is never globally decreasing for periods longer than K periods. This is for example the case from 1972 to 2023 with K equal to 96 months or more: the two methods yield *exactly* the same result.

As discussed above, and this was our main reason to depart from the isotonic regression, we would like our method to automatically be able to track long-term increases, but also decreases, in the temperature. Therefore, it is also useful to define similar conditions and values to determine when the temperature goes below a certain threshold L from period i (that will hopefully become useful later in this century):

$$\bar{T}_i^K = \max_{p\in[i,i+K-1]} T_{i,p} \geq L. \tag{6}$$

Here, $\bar{T}_i^K$ is the maximum temperature that will be observed on average from period i to any future period, within a horizon of K periods. Observe that by construction, we always have that $\underline{T}i^K \leq T_i \leq \bar{T}_i^K$.

We could stop here, and, to the threshold question, we could simply answer that it is the earliest period i for which $\underline{T}_i^K$ is at least L (and similarly to be below the threshold L). However, for the sake of clarity and simplicity, it would be helpful to be able to plot the evolution of the interesting portions $\underline{T}i^K$ and $\bar{T}_i^K$ over time as a single time series. To that end, we recursively define the following time series $\tilde{T}_i^K$, that we will call the K-minmax average, or simply the minmax average:

$$\tilde{T}_i^K = \begin{cases} T_i & \text{if } i = M, \\ \tilde{T}_{i-1}^K & \text{if } \underline{T}_i^K \leq \tilde{T}_{i-1}^K \leq \bar{T}_i^K, \\ \underline{T}_i^K & \text{if } \tilde{T}_{i-1}^K < \underline{T}_i^K, \\ \bar{T}_i^K & \text{if } \tilde{T}_{i-1}^K > \bar{T}_i^K. \end{cases} \tag{7}$$

The minmax average $\tilde{T}$ starts as equal to T in $i = M$, then is kept constant, unless it does not fall in the range $[\underline{T}_i, \bar{T}_i]$, in which case it is brought within that interval. Broadly speaking, $\tilde{T}_i^K$ will track the highest past value of $\underline{T}i^K$ while the temperature T_i is globally increasing, while it will conversely track the lowest past value of $\bar{T}_i^K$ if T_i is globally decreasing. This makes it possible to directly read from the unique series $\tilde{T}_i^K$ when we pass thresholds in both directions.

Applications of the minmax average methodology

1850–2023 Dataset

We first apply our methodology to the monthly GMSTA time series, for $K = 120$ months $= 10$ years and $K = 240$ months $= 20$ years in Fig. 4.

Strictly speaking, the values from 2004 cannot be determined for $K = 20$ and from 2014 for $K = 10$, but here we just assumed that the temperature will not go back under $+1.29$°C on average from any time interval starting in July 2023 and ending before 2044, which is a fairly safe bet today [5].

As expected, the minmax average is able to track the medium-to-long term ups and downs in the data, while smoothing out the shorter term oscillations. Choosing $K = 10$ years

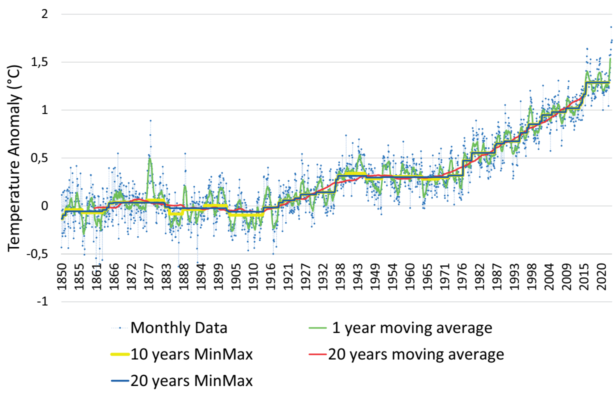

Figure 4. Temperature Anomaly from period 1850–1900, 1- and 20-years moving averages and 10- and 20-years minmax averages, Berkeley Earth Data [8].

Table 1. Range of uncertainties expressed as 95% confidence intervals in the Berkeley Earth Temperature Record dataset [8]

	1972–2022	1850–2022, up to
Monthly avg	0.030–0060	0.5
1-year avg	0.025–0040	0.2
5-years avg	0.021–0025	0.2
10-years avg	0.020–0022	0.12
20-years avg	0.017–0021	0.09

naturally yields a time series that more dynamically tracks the ups and downs in the data compared to $K = 20$. But on balance, using $K = 20$ years seems a better option, first of all because it is in line with the current usage of IPCC to allow averaging out within intervals of 20 years, and second because the behavior of the time series 10 years-minmax seems a bit too dynamic especially in the period 1870–1910. In particular both the 20-years moving average and the 20-years minmax average are monotonously decreasing in that period, while the 10-years minmax average increases between 1885 and 1900. But we want to highlight the fact that this is of little importance in the current regime of rapid temperature increases, and that even taking $K = 8$ years would not make any difference since 1972. Rather, this will become more relevant in the future, when the temperatures will stabilize, and the discussion will revolve around the question "Is the global temperature actually decreasing?".

Let us turn now to the question of the uncertainty around the minmax values computed. Since the minmax average only depends on the time series T, the uncertainty about the minmax average only depends as well on the uncertainty of this underlying data, and in particular in the joint distribution of these. In practice, we are interested in applying the methodology to the GMSTA. The Berkeley Earth dataset we use report uncertainties in terms of 95% confidence intervals, for monthly, 1-year, 5-years, 10-years and 20-years averages [8]. Table 1 summarizes typical ranges of these values.

Without knowledge about the joint distributions of the errors in different time periods, we cannot estimate similar uncertainties for the minmax average, but we can still give reasonably tight bounds from these as follows.

By construction, the minmax average at a certain period is the average of the GMSTA within an interval (usually the constant interval to which the period belongs to). This immediately implies that the uncertainty of the minmax average will be at most the monthly uncertainty in the GMSTA, and usually lower, since it is computed as the average over a certain number of time periods. Also, as with the simple multi-years averages, we can expect the uncertainty of the minmax average to be smaller when it corresponds to the average of the temperature over longer intervals.

Now observe that the intervals of constant minmax average are between 1 and 125 months, and usually several years in duration. We can therefore conclude that the uncertainty of the minmax average (expressed as a 95% confidence interval) in the recent period 1972–2022 will usually be of the order of 0.025–0030 or lower, which is small. Therefore, in the remaining of this study, we will only analyze the central case.

Synthetic time series

Applying the method to generated data allows to better compare the minmax average to the processes underlying the data. We first compute the minmax average for the five main IPCC scenarios up to 2100 (Fig. 5). For SSP2-4.5 and higher, the minmax average is simply equal to the time series itself, as they are monotone. For SSP1-1.9 and SSP1-2.6, the top of the peak is simply truncated, but by less than 0.01°C because the maximum is nearly flat in both cases.

A more insightful experiment is to generate time series as a smooth curve with a maximum (similar to SSP1 scenarios above) plus an autoregressive process with mean zero. We generated a hundred such time series, computed the minmax average for each, and report relevant statistics in Fig. 6, together with an illustrative single realization of the stochastic process. We can conclude the following. The median minmax average correctly identifies the trend, or, more precisely, the minmax average of the trend (i.e. trend with the peak truncated). This would justify using the minmax average of the median scenario as a correct tool to determine GMSTA thresholds attained at some future date with 50% probability. Moreover, the median minmax average converges much faster to the minmax average of the trend compared to the rate of convergence of the median of the underlying data (not plotted for clarity) to the trend. A similar observation is that the 5^{th} and 95^{th} percentiles of the minmax averages are closer to trend than the corresponding percentiles of the underlying data. This behaviour is expected as the minmax average implicitly averages over time, reducing the variance of the error with mean zero.

The single realization illustrated is also of interest. For the sake of the discussion, let us assume the stochastic process represents our best (exact in this case) climate model, and the realization the historical temperature as revealed over time. The minmax average is first increasing, reaches a maximum, then is decreasing, matching the overall trend of the stochastic climate process. But there are substantial differences as well. In the realization selected, the maximum minmax is reached well before the maximum of the trend (2045 versus 2054) and is also lower in magnitude (2.55°C versus 2.62°C). Also the minmax average of the single realization has a much less smooth rate of increase/decrease than the trend. Simply put, the stochastic error might cause an actual temperature record that is qualitatively different from the trend. This is also consistent with the actual practice by IPCC, where sentences of the type "there is a 50% probability that the temperature will stay below 2°C if the cumulative emissions from 2023 stay below 1150Gt" are common, acknowledging that the actual temperature record could be substantially and in many respects different from the central case [6].

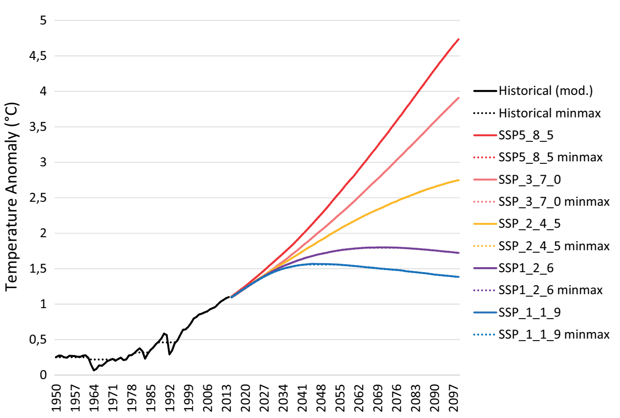

Figure 5. Main IPCC scenarios and associated minmax averages.

Analysis of the recent period 1972–2022

In this section, we want to have a closer look to the temperature record in the last 50 years, its relationship to the 20 years minmax average and then Oceanic Niño Index [27]. We will discuss in the next Section the recent high temperatures of 2023 and what this might mean for the minmax average today.

We can see in Fig. 7 that we passed the 0.5°C threshold in 09/1979, and were already close to it at 0.47°C as soon as 12/1976, just after a fairly large one-month increase of +0.154°C. From

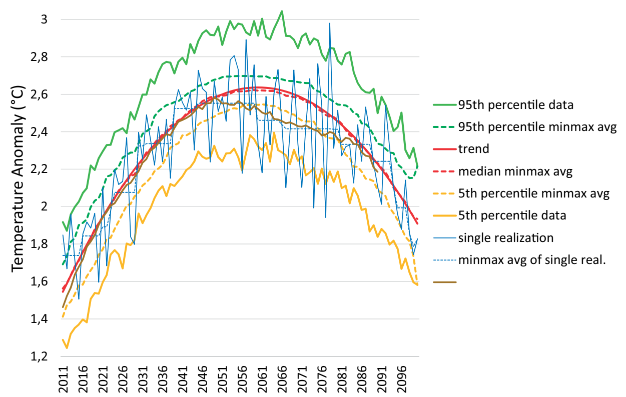

Figure 6. Minmax averages of a smooth peaking trend plus zero-mean autoregressive stochastic process, and associated statistics.

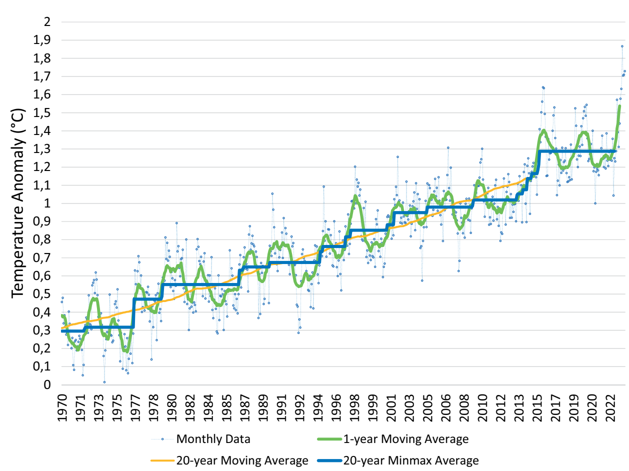

Figure 7. Temperature anomaly from period 1850–1900, 1- and 20-years moving average and minmax average for the period 1970–2023, Berkeley Earth Data [8].

1976 to 2013, the increase is much more gradual: the maximum increase from one month to the next is +0.066°C in 2001, and these increases are fairly evenly spaced. The threshold of +1°C is passed in 06/2009, although the minmax average stays very close to +1°C, between 0.95 and 1.02, from 11/2001 to 07/2013. This corresponds to the infamous "hiatus" widely discussed in the media towards the end of that 12-years stretch [28].

Then the minmax average increases rapidly from 1.02°C in 07/2013 to 1.29°C in 10/2015, a + 0.27°C increase in just 27 months. This corresponds to the fairly strong El Niño phase of 2014–2016. Note that, by construction, the minmax value of 1.29°C is equal to the average temperature anomaly in the period 10/2015 to 01/2023. Over the period 1972 to 2019, a period long enough to smooth out the ENSO cycle, most approaches coalesce on a rate of increase of roughly +0.165°C per decade on average [8, 16].

One striking feature when looking at Fig. 7 is that to each constant interval of the minmax average typically corresponds one or two up-and-down oscillations of the 1-year moving average, starting first above the constant segment. For example, on the flat segment from 09/1997 to 02/2001, the 1-year moving average is larger than the minmax average from 09/1997 until 12/1998, then below until 02/2001. Of course this is nearly by construction, since we know property (e) holds within each such interval. But this seems to corresponds to the ups and downs of the ENSO oscillation as observed in Fig. 8, where two time series are plotted. The first one is the ONI lagged by two months and the other one is the difference between a 3-months moving average of the temperature anomaly minus the minmax average. (The value of 3 months has been chosen for two reasons. The first is that the ONI index is itself computed by averaging over 3 months. The second one is that the minmax average does not anticipate rapid increases in the temperature anomaly, while a 12-months moving average does, so taking the difference between these two values would compare two qualitatively different indicators.) One can clearly see a general correspondence between the timing and the magnitude of the extrema of the two curves. This is not surprising since it is well known that the ENSO cycle is correlated with the global temperature, and that the minmax average $\tilde{T}_i$ is by construction the average temperature anomaly over a certain time interval containing i.

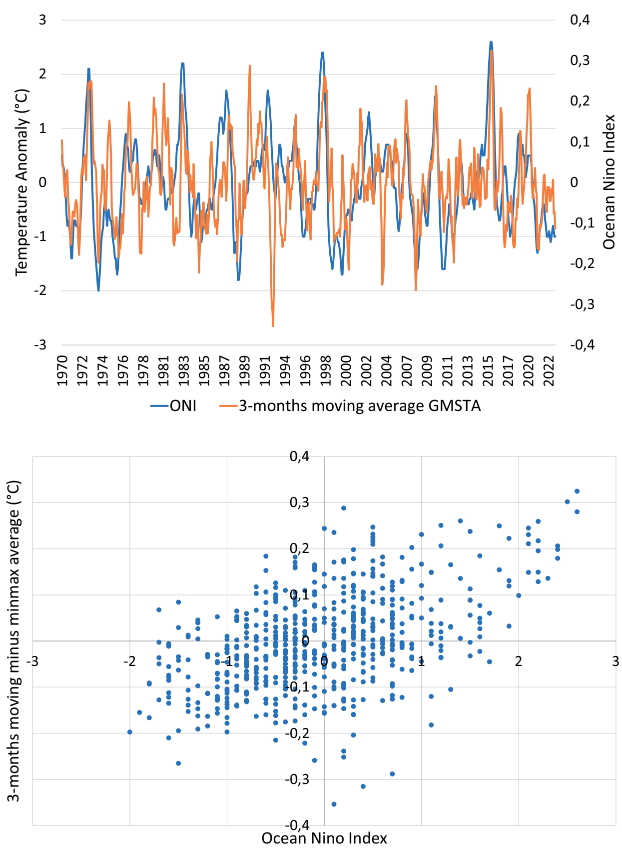

Figure 8. 2-months lagged Oceanic Niño Index and the difference between the 3-months average and the minmax average: Time series and Scatter plot, Berkeley Earth Data and NOAA [8, 27].

There are a few outliers, notably the 3-months spike in temperature in the period 3–5/1990 that corresponds to a neutral ONI. But there is a general correlation between these two time series as shown in the corresponding scatter plot. This demonstrates that ENSO is a major driver of the difference between the temperature anomaly and its minmax average.

Outlook for 2023

The temperature anomalies in 2023 were suddenly much higher, especially in the second half of the year, in line with an El Niño event starting in May. The average GMSTA for 2023 is 1.54°C, with a monthly maximum of 1.87 in September, while the last 6 months average at 1.70°C. The main question if of course: what will the minmax average end up being mid-2023? With a choice of K corresponding to 20 years, strictly speaking, we will need to wait until 2043 to answer this question. However, since 1972, using values for K as low as 8 years makes no difference. Furthermore, over the same period, the maximum of $\tilde{T}_i^4 - \tilde{T}_i^{20}$ ($K = 4$) is 0.032, while the maximum of $\tilde{T}_i^3 - \tilde{T}_i^{20}$ ($K = 3$) is 0.044. Therefore, we will have a good idea already in mid-2026 or mid-2027, or in general after the end of the next La Niña phase of ENSO. This is because we know that the current warmer phase will be followed by a cooler one, and our minmax average measure is essentially based on averaging these, starting from June or July 2023, when the temperature anomaly reached 1.44 and 1.58 respectively.

To support policy effectively, we would want to obtain an estimate today, instead of waiting. To that end, as discussed above and in Betts *et al.* we would need to simulate the global climate for the next three to four years at least, with a sufficient precision concerning the duration and the magnitude of the ENSO cycle, both up and down [9]. Forecasting the ENSO cycle is currently a very active area of research, but the maximum horizon of current models is 22 months, with most models limited to 12 months [29, 30]. Moreover, La Niña phase lasts typically longer and is substantially more difficult to predict than the Niño phase [31]. But forecasting the duration and depth of the cold part of the cycle is actually key to estimate the minmax average today, in the middle of the warm phase. Therefore, we believe that more research is necessary for using climate simulation models to meaningfully estimate the minmax average today.

In the meantime, some simple static measures can still help shed light on the question. A first approach, based exclusively on data up to mid-2023, is to use an average rate of increase per year. Up to now, the main difficulty was to determine meaningful starting and end time points to compute this average. Indeed, because of the ENSO cycle, selecting for example a high point in the cycle as a starting point and a low one as an end point would result in an artificially low average rate of increase. But selecting two high (or low) points is also problematic because some ENSO cycles are much stronger than others, so that a single monthly or yearly outliers could have an outsize impact on the estimate, rendering the obtained value somewhat arbitrary. A good solution to these issues have been to build such averages over long time intervals, resulting for example in the +0.165°C per decade for the period 1972–2019 discussed above. But this makes it impossible to closely track a variation in time in this rate of increase.

The minmax average approach offers a simple solution to these issues, by dynamically identifying time intervals when the global temperature is constant and possibly spanning multiple ENSO cycles, and short periods of relatively rapid increases that are not reversed later on average. Firstly, it is now easy to select points in time similarly placed in this rapid-increase/constant-level cycle. Secondly, the minmax average values in a constant interval are themselves averages over several years, and so relatively insensitive to outliers.

Table 2 gives the rate of increase in the minmax average GMSTA in intervals starting just after periods of rapid increase in the minmax average GMSTA, or equivalently at the start of intervals of constant minmax average GMSTA. (Note that the increase in 12/1976 happened after a very long period of nearly constant minmax average since 1937.) In the period from 12/1976 to 01/1995, the rate of increase is on average 0.161°C per decade, while it is 0.253°C from 01/1995 to 10/2015, a clear acceleration. In that perspective, the rapid increase in GMSTA in the period 2013–2015 can simply be seen as catching up on a trend that was already present in 1995–2001, after the hiatus of 2001–2013.

Now, assuming (i) this trend since 1995 continues, and (ii) a new rapid increase in the minmax average in 2023 (as now actually observed) will merely catch up on the trend after the hiatus of 2015–2022, we would expect a minmax average of $1.29 + 0.025*8$ or 1.49°C in 2023.

Let us now try to make use of the actual temperature observed in the second half of 2023 for our purposes. We already observed that the minmax average, during an interval of constant value, is typically the mean of the anomaly during one warm-cold cycle (and more rarely two), starting with the warm phase. Suppose we can estimate (i) the mean temperature over the warm phase of the cycle and (ii) by how much the cold phase of the cycle will later bring down the average over the full cycle, we can estimate the minmax average by just taking the difference between these two values.

Now, we can in early 2024 only observe the temperature anomaly for the initial months of the warm phase, which is not currently over. So what we can do is to apply the methodology outlined in the previous paragraph, but using more short-term averages than the mean temperature over the full warm phase. In particular, Table 3 gives the maximum differences observed between the k-consecutive-months averages and the minmax average, across all the months in the 1970–2022 period.

Table 2. Average rate of increase in the minmax average GMSTA in intervals starting just after periods of rapid increase in the minmax average GMSTA

Time	Minmax avg	Avg rate of incr. per decade
déc-76	0.472	
sept-79	0.553	0.294
juin-87	0.650	0.125
janv-95	0.762	0.148
sept-97	0.852	0.336
nov-01	0.949	0.233
oct-15	1.288	0.243

Table 3. Maximum differences observed between the k-consecutive-months averages and the minmax average, across all the months in the 1970–2022 period

k	Timing	Average	Minmax	Difference
1	3/1990	1.05	0.67	0.38
3	1–3/2016	1.61	1.29	0.32
6	12/1972–5/1973	0.57	0.32	0.25 (tie w 1998,2016)
12	9/1997–8/1998	1.04	0.85	0.19

Assuming these record differences between these short term and minmax averages will not be broken this ENSO cycle starting in 2023, the implied lower bounds for the minmax average today are 1.49 for $k = 1$ (anomaly of 1.87 in 09/2023 minus 0.38), 1.44 for $k = 3$ (average anomaly of 1.76 in 9–11/2023 minus 0.32) and 1.45 for $k = 6$ (anomaly of 1.70 in 07–12/2023 minus 0.25). The highest recent 12-months anomaly is 1.54 for the full year 2023, but usually the maximum is attained later in the El Niño phase, typically for an interval centered in the month of January or February, so this is not too informative yet.

Moreover, according to Fig. 8, record differences between minmax and short-term averages are typically attained for higher (2-months lagged) ONI values. The maximum for the 3-months average was attained for the month of October 2023, and the August 2023 (2-months lagged) ONI was 1.3. Therefore using a record difference of 0.26 for $k = 3$ (instead of 0.32), typical for this range of ONI values as illustrated in the scatter plot of Fig. 8, seems justified, resulting in a stronger lower bound of 1.5.

There are other forcings that might explain the unusually high 2023 temperature increase: the solar cycle that is in an ascending phase, the water vapor ejected in the atmosphere by the Hunga Tonga eruption, or the recent reduction in marine fuel pollution [10, 32, 33]. However, these forcings, smaller in magnitude than the effect of greenhouse gases and the ENSO cycle, are expected to stay at similar levels for the next 3–4 years, and therefore cannot be expected to compensate the high temperatures of 2023–2024 later during this ENSO cycle.

It is good of keep in mind that (i) the current El Niño phase is not over so that the max k-averages values for this 23–24 cycle might even be higher, and (ii) these estimated lower bounds are derived from record historical differences. Therefore, for the current minmax average to be well below 1.5°C, for example below 1.45°C, all these record differences must be exceeded. That would imply much stronger and/or longer cold phase in the ENSO cycle than observed historically to compensate for the record global temperature anomalies measured in 2023. Rather, it seems that a typical cold phase in 2024–2026 would result in a minmax average temperature anomaly in 2023 of the order of 1.5°C, or more. We want to emphasize that this is not an estimate, nor the lower end of a confidence interval. The goal of this analysis is simpler: to show that such a result is entirely plausible, without attaching a probability to it. As discussed, this would require better simulation models over the full ENSO cycle.

Discussion

We summarize here the advantages of the proposed approach.

- The Paris Agreement (and its ultimate goal, the well-being of human societies) is about the actual temperature anomaly, not about the man-made, non-cyclic or non-stochastic fraction of it (second section).

- The proposed methodology makes statistical sense, and is coherent with the objective of averaging out temporary increase or decrease of the temperature that are reversed in the short-to-medium term (third and fourth sections).
- The minmax average tracks the data exactly when it is monotone, and as close as possible otherwise (third section).
- The methodology is simple to explain and justify (e.g. to non-scientists), and minimizes assumptions made. It is therefore more likely to be widely accepted (fourth section).

It seems that the main issue with the new approach will be that the minmax average time series will typically exhibit “alternating plateaus and jumps”, or, more formally, that the evolution of its *rate of increase* is not smooth enough over time. Note that we could easily modify the optimization problem (1)–(2) to obtain such a smoother solution. But why should we do this? Firstly, in abstract, it is hard to see why the evolution of the rate of increase is important for answering the threshold question, which is at heart a question regarding the level. This seems a category error, similar to confounding a function and its second derivative. Second, doing so would create the issues discussed in the second section. Also, as already pointed out, there is nothing in the proposed methodology that imposes a non-smooth behavior, it merely leaves it possible. For example, when applied to the SSP1 scenarios in Fig. 5, the obtained minmax series is smooth. If there are plateaus and jumps when applied to GMSTA, it is because these are implicitly present in the temperature record, as the discussion about the “hiatus” of 2001–2013 illustrated. Now, it is clear that the human contribution to the GMSTA is expected to evolve smoothly over time, as the stock of GHG in the atmosphere itself evolves slowly over time. Therefore, insisting that to answer the threshold question one must use a first-order-smoothed GMSTA is a way to implicitly constrain the methodology to mainly depend on its man-made contribution. Once this constraint is removed, it is hard to see why a smooth rate of increase is important. Finally, it is true that a smooth rate of increase would enable more accurate forecasts of the date a threshold is passed. Sudden variations in the temperature record, if they are substantial in magnitude, and their timing is hard to predict, will increase the uncertainty of these forecasts, and this needs to be reflected in a central scenario and comparatively large uncertainty estimates. But, as shown in the synthetic data section, even using the minmax average, this results in smooth median and quantile time series. It is only once the forecasts become measurements that we might observe a non-smooth minmax average.

Summary and conclusion

After discussing why current approaches to answer the threshold question are not fully satisfactory, we have described a new approach (the minmax average) that is agnostic about attribution. We have outlined its close relationships with the standard monotone (or isotonic) regression. We also discussed its advantages and tried to answer expected critiques. When applied to the GMSTA record, we obtain the following conclusions. First, the rate of increase in the temperature anomaly has been $\sim$+0.25°C per decade since 1995. Second, based on this new estimate alone and assuming the minmax average will increase in 2023, we should plausibly expect the minmax average of the temperature anomaly to reach 1.49°C in 2023. Factoring in the record-shattering actual temperatures of the second half of 2023, we conclude that having breached the 1.5°C already in July 2023 seems entirely plausible. This would be much earlier than the central estimate of 2030–2032 of the IPCC [6, 9]. Moreover, if the current trend of +0.25°C per decade is maintained, we can expect to reach the threshold of +2°C in the first El Niño phase after 2043 (so in 2043–2050), or even earlier if an acceleration of the rate of increase materializes [34, 35]. On the contrary, if vigorous emission reductions are achieved, the rate of temperature increase could slow down rapidly [36].

We will nonetheless have to wait until the end of the next cold phase of the current ENSO cycle, probably in 2026 or 2027, to be able to meaningfully estimate the new minmax average GMSTA reached in July 2023. Improvements in forecasting early the mean GMST in a full ENSO cycle would results in an earlier such estimate.

Of course, based on the way the indicator is constructed, and on the past history, we can expect the minmax average GMSTA to stay constant at its new value attained in 2023 for several years, be it below or above 1.5°C. Even the former case should not lead to complacency. This will likely be very close to the agreed limit of 1.5°C already.

Open research section

The temperature anomaly compiled by the Berkeley Earth group can be accessed at https://berkeley-earth-temperature.s3.us-west-1.amazonaws.com/Global/Land_and_Ocean_complete.txt. The ONI index from NOAA National Weather Service, Climate Prediction Center is accessible at https://origin.cpc.ncep.noaa.gov/products/analysis_monitoring/ensostuff/ONI_v5.php.

Acknowledgements

This work was carried out while on sabbatical leave at the Sauder School of Business, University of British Columbia. I am grateful for comments from François Massonnet on a very early draft version of this study. I also want to thank the two anonymous referees for their fruitful comments.

Author contributions

Mathieu Van Vyve (Conceptualization [equal], Data curation [equal], Formal analysis [equal], Funding acquisition [equal], Investigation [equal], Methodology [equal], Project administration [equal], Resources [equal], Software [equal], Supervision [equal], Validation [equal], Visualization [equal], Writing—original draft [equal], Writing—review & editing [equal])

Conflict of interest

None declared.

Funding

No specific funding was used to carry out this research.

References

1. Kornhuber K, Lesk C, Schleussner CF *et al*. Risks of synchronized low yields are underestimated in climate and crop model projections. *Nat Commun* 2023;**14**:3528.

2. Khojasteh D, Haghani M, Nicholls RJ *et al.* The evolving landscape of sea-level rise science from 1990 to 2021. *Commun Earth Environ* 2023;**4**:257.
3. Liu W, Lim WH, Sun F *et al.* Global freshwater availability below normal conditions and population impact under 1.5 and 2°C stabilization scenarios. *Geophys Res Lett* 2018;**45**:9803–13.
4. Lenton TM, Xu C, Abrams JF *et al.* Quantifying the human cost of global warming. *Nat Sustain* 2023;**6**:1237–47.
5. Matthews HD, Wynes S. Current global efforts are insufficient to limit warming to 1.5°C. *Science* 2022;**376**:1404–9.
6. Masson-Delmotte V, Zhai P, Pirani A *et al.* (eds). Climate change 2021: the physical science basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, 2023. Cambridge University Press, Cambridge, NY, USA.
7. McCulloch MT, Winter A, Sherman CE, Trotter JA. 300 years of sclerosponge thermometry shows global warming has exceeded 1.5°C. *Nat Clim Chang* 2024;**14**:171–7.
8. Rohde RA, Hausfather Z. The Berkeley earth land/ocean temperature record. *Earth Syst Sci Data* 2020;**12**:3469–79.
9. Betts RA, Belcher SE, Hermanson L *et al.* Approaching 1.5°C: how will we know we have reached this crucial warming mark? *Nature* 2023;**624**:33–5.
10. Diamond MS. Detection of large-scale cloud microphysical changes within a major shipping corridor after implementation of the international maritime organization 2020 fuel sulfur regulations. *Atmos Chem Phys* 2023;**23**:8259–69.
11. Savitzky A, Golay MJE. Smoothing and differentiation of data by simplified least squares procedures. *Anal Chem* 1964; **36**:1627–39.
12. Hasselmann K. Multi-pattern fingerprint method for detection and attribution of climate change. *Climate Dynamics* 1997; **13**:601–11.
13. Stott PA, Gillett NP, Hegerl GC *et al.* Detection and attribution of climate change: a regional perspective. *WIREs Clim Chang* 2010; **1**:192–211.
14. Bindoff, N.L., Stott, P.A., AchutaRao, K.M., Allen, M.R., Gillett, N., Gutzler, D., Hansingo, K., Hegerl, G., Hu, Y., Jain, S., Mokhov, I.I., Overland, J., Perlwitz, J., Sebbari, R. and Zhang, X.. *Detection and Attribution of Climate Change: From Global to Regional*. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, NY, USA, 2013, p. 867–952.
15. Zhai P, Zhou B, Chen Y. A review of climate change attribution studies. *J Meteorol Res* 2018;**32**:671–92.
16. Haustein K, Allen MR, Forster PM *et al.* A real-time global warming index. *Sci Rep* 2017;**7**:15417.
17. Diffenbaugh NS, Barnes EA. Data-driven predictions of the time remaining until critical global warming thresholds are reached. *Proc Natl Acad Sci U S A* 2023;**120**:e2207183120.
18. Kruskal JB. Nonmetric multidimensional scaling: a numerical method. *Psychometrika* 1964;**29**:115–29.
19. Fielding A, Barlow RE, Bartholomew DJ *et al.* Statistical inference under order restrictions. the theory and application of isotonic regression. *J R Stat Soc Ser A* 1974;**137**:92–3.
20. Lee C-IC. The min-max algorithm and isotonic regression. *Ann Statist* 1983;**11**:467–77.
21. Best MJ, Chakravarti N. Active set algorithms for isotonic regression; a unifying framework. *Math Prog* 1990;**47**:425–39.
22. Oron AP, Flournoy N. Centered isotonic regression: Point and interval estimation for dose-response studies. *Stat Biopharm Res* 2017;**9**:258–67.
23. Luss R, Rosset S, Shahar M. Efficient regularized isotonic regression with application to gene-gene interaction search. *Ann Appl Stat* 2012;**6**:253–83.
24. Niculescu-Mizil A, Caruana R. Predicting good probabilities with supervised learning. In: *Proceedings of the 22nd International Conference on Machine Learning, New York, NY, USA, ICML '05*, p. 625–32.
25. Kalai A, Sastry R. The isotron algorithm: high-dimensional isotonic regression. In: *COLT 2009 - The 22nd Conference on Learning Theory, Montreal, Quebec, Canada*, 2009.
26. Ghazi B, Kamath P, Kumar R, Manurangsi P. Private isotonic regression. In: Koyejo S, Mohamed S, Agarwal A, Belgrave D, Cho K, Oh A (eds), *Advances in Neural Information Processing Systems*, vol. **35**. Curran Associates, Inc., 2022, pp. 8577–90.
27. Huang B, Thorne PW, Banzon VF *et al.* Extended reconstructed sea surface temperature, version 5 (ersstv5): upgrades, validations, and intercomparisons. *J Clim* 2017;**30**:8179–205.
28. Medhaug I, Stolpe MB, Fischer EM, Knutti R. Reconciling controversies about the 'global warming hiatus'. *Nature* 2017;**545**:41–7.
29. Wang H, Hu S, Li X. An interpretable deep learning enso forecasting model. *Ocean-Land-Atmos Res* 2023;**2**:0012.
30. Rivera Tello GA, Takahashi K, Karamperidou C. Explained predictions of strong eastern pacific el niño events using deep learning. *Sci Rep* 2023;**13**:21150.
31. Chen H, Jin Y, Shen X *et al.* El niño and la niña asymmetry in short-term predictability on springtime initial condition. *NPJ Clim Atmos Sci* 2023;**6**:121.Aug.
32. Gray LJ, Beer J, Geller M *et al.* Solar influences on climate. *Rev Geophys* 2010;**48**:RG4001.
33. Schoeberl MR, Wang Y, Ueyama R *et al.* The estimated climate impact of the hunga Tonga-hunga ha'apai eruption plume. *Geophys Res Lett* 2023;**50**:e2023GL104634.
34. Richardson MT. Prospects for detecting accelerated global warming. *Geophys Res Lett* 2022;**49**:e2021GL095782.
35. Hansen JE, Sato M, Simons L *et al.* Global warming in the pipeline. *Oxford Open Clim. Change* 2023;**3**:kgad008.
36. Dvorak MT, Armour KC, Frierson DMW *et al.* Estimating the timing of geophysical commitment to 1.5 and 2.0°C of global warming. *Nat Clim Chang* 2022;**12**:547–52.

Oxford Open Climate Change, 2024, 4, 1–10
https://doi.org/10.1093/oxfclm/kgae009
Research Article